Practical Electronics

The UK's premier electronics, computing and maker magazine

Max's Cool Beans
Cellular automata: The Game of Life

Net Work
Uses for 18V/21V Li-ion tool batteries

Audio Out
RetroTech report part 2

Circuit Surgery
Accurately simulating 555s with LTspice

DIY DCC Model Railway Part 1: DCC Decoder

Ducted Heat Transfer Controller

Rigol DHO924S Digital Storage Oscilloscope Review

PLUS!

Techno Talk – new web development & testing techniques

The Fox Report – Humanoid robots are just around the corner

Advanced uses for HomeAssistant running on a Raspberry Pi

Oct 2026 £5.99

9 772632 573030 10

PCBWay

Beyond Standard —

Discover PCBWay's Advanced PCB Lineup

About PCBWay:

With dedicated workshops for advanced PCBs, PCBWay offers full-spec, high-precision, and large-scale production. We ensure customers are supplied with PCBs that meet stringent quality standards at the lowest cost.

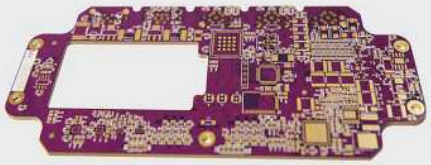

HDI PCB

- 4-60 layers
- FR-4, High Tg, Halogen-free, PTFE
- Finished Copper: 1/3OZ–8OZ
- Min Track/Spacing (inner): 2/2mil (H/HOZ base copper)
- Min BGA Pitch: 0.35mm

High Frequency PCB

- RO4003C, RO4350B, Ro3003, Ro3010, RT5880
- Min. Track/Spacing: 2mil
- Min. Hole Size: 0.15mm
- Finished Cooper: 1-2oz
- Board Thickness: 0.2–3.2mm

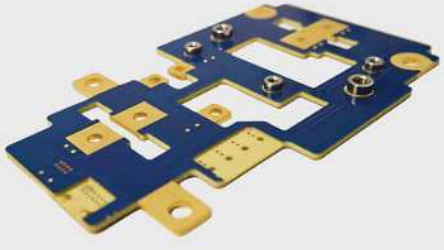

Other PCBs with Special Process

- Embedded Copper PCB
- Backplane PCB
- Cavity PCB
- White/Partially Hybrid PCB
- Black Core PCB (After-dark)

www.pcbway.com service@pcbway.com

Practical Electronics

Volume 55. No. 10
October 2026
ISSN 2632 573X

Contents

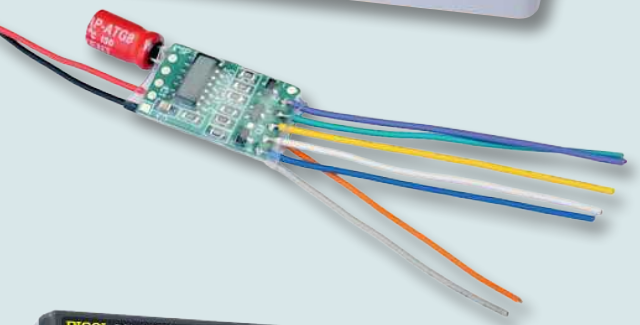

Projects and Circuits

DCC Decoder for model locomotives by Tim Blythman **14**
This DCC Decoder project provides the first essential part of a DIY DCC system. We will follow up with many more DCC components over following issues, including a Base Station, Booster, Remote Controller, Destination Display and much more.

Ducted Heat Transfer Controller, part one by Julian Edgar and John Clarke **24**
This smart controller can improve the energy efficiency of your home. It transfers warm or cool air between rooms automatically when needed, using wall-mounted temperature sensors hidden behind plates and a small control panel.

Series, Features and Columns

Buying Second-Hand Speakers by Julian Edgar **4**
If you're shopping for second-hand audio gear, bargains can be had, but you need to know what you're looking for.

Power Electronics, part seven by Andrew Levido **8**
Higher switching frequencies can make input and output filters simpler, with smaller magnetics. They also allow a faster response to changes in the load current. Switching losses become a major problem at higher frequencies, but there is a solution.

Rigol RHO924S DSO & DM858 bench DMM review by Tim Blythman **32**
These recent releases from Rigol are affordable and also very portable. Despite that, they have very creditable performance and many useful features.

Net Work by Alan Winstanley **40**
Alan's occasional column delves into popular 21V rechargeable tools and batteries. He brings updates on the forthcoming 'Temu Tax' and ruminates about buying an EV.

Using Home Assistant, part two by Richard Palmer **44**
Having described how to set up a custom Home Automation (HA) system using a Raspberry Pi last month, we now look at features like cameras, custom dashboards, IR remote control, notifications and remote access on a mobile phone, tablet or PC.

Audio Out by Jake Rothman **50**
The RetroTech UK event and radio rallies, part two of two. RetroTech UK is run by the British Vintage Wireless Society (BVWS). It's a wonderful sale of old electronic stuff that takes place every May. I saw many interesting things when I attended this year.

Circuit Surgery by Ian Bell **56**
We continue our examination of the popular 555 timer IC from last month. This time we consider how to create realistic simulations of circuits that include the 555.

The Fox Report by Barry Fox **64**
It won't be long before you can buy your very own robot.

Techno Talk by Max the Magnificent **66**
Web development has changed a lot in the last few years!

Max's Cool Beans by Max The Magnificent **68**
Weird & Wonderful Arduino Projects, Part 22: It's life, Jim, but not as we know it! Let's implement celluar automata, including Conway's Game of Life, Seeds and Brian's Brain. We also work on the 7-segment VFD-based Chronological Contraption.

Regulars and Services

Editorial **2**
Clear and simple mathematical notation

Reader services – Editorial and Advertising Departments **2**

Subscribe to Practical Electronics and save money **3**

Practical Electronics PCB Service **77**
PCBs for Practical Electronics projects

Classified ads and Advertiser index **79**

Next month! – highlights of our next issue of *Practical Electronics* **80**

Made in Australia & the UK.

Written in Britain, Australia, the USA and Ireland.

Read everywhere.

© Silicon Chip Publications Pty Ltd 2026

Copyright in all drawings, photographs, articles, technical designs, software and intellectual property published in *Practical Electronics* is fully protected, and reproduction or imitation in whole or in part are expressly forbidden.

The November 2026 issue of *Practical Electronics* will be published on Tuesday, 6th of October, 2026 (see p80).

Renewals paid by PayPal, Visa and Mastercard are now handled via our website: www.electronpublishing.com

Practical Electronics

Volume 55. No. 10
October 2026
ISSN 2632 573X

Editorial

Editorial offices
Electron Publishing (Australia)
Tel +61 2 9939 3295
Email pe@pemag.au
Web www.electronpublishing.com

Address mail to:
Electron Publishing (Australia)
PO Box 194, Matraville NSW 2036
Australia

Advertising enquiries
+61 2 9939 3295 pe@pemag.au

Editor	Nicholas Vinen
Publisher	Nicholas Vinen
Digital subscriptions	Stewart Kearn **Tel 07918 614662**
Online Editor	Alan Winstanley
Web Systems	Kris Thain
Production	Bao Smith
Technical staff	Tim Blythman, John Clarke

Print subscriptions
Practical Electronics Subscriptions
PO Box 6337
Bournemouth BH1 9EH **Tel** 01202 087631
United Kingdom **Email** pesubs@selectps.com

Technical enquiries
We regret that technical enquiries cannot be answered over the telephone. We are unable to offer any advice on the use, purchase, repair or modification of commercial equipment or the incorporation or modification of designs published in the magazine.

Questions about articles or projects should be sent to the editor by email: pe@pemag.au

Projects and circuits
All reasonable precautions are taken to ensure that the advice and data given to readers is reliable. We cannot, however, guarantee it and we cannot accept legal responsibility for it.

Some projects and circuits published in *Practical Electronics* employ voltages that can be lethal. Do not build, test, modify or fix any mains-powered equipment unless you fully understand the safety aspects involved and you use an RCD (GFCI) adaptor.

Component supplies
Silicon Chip Publications may offer kits or other parts for making our projects, but not in all cases. When kits are not available, readers will need to find and source parts themselves. We advise readers to check that all parts are still available before commencing any project in a back-dated issue.

Advertisements
Although the proprietors and staff of *Practical Electronics* take reasonable precautions to protect the interests of readers by ensuring as far as practicable that advertisements are bona fide, the magazine and its publishers cannot give any undertakings in respect of statements or claims made by advertisers, whether these advertisements are printed as part of the magazine, or in inserts. The Publishers regret that under no circumstances will the magazine accept liability for non-receipt of goods ordered, or for late delivery, or for faults in manufacture.

Transmitters/bugs/telephone equipment
We advise readers that certain items of radio transmitting and telephone equipment which may be advertised in our pages cannot be legally used in the UK. Readers should check the law before buying any transmitting or telephone equipment, as a fine, confiscation of equipment and/or imprisonment can result from illegal use or ownership. The laws vary from country to country; readers should check local laws.

Clear and simple mathematical notation

Sometimes interesting I noticed about serious mathematicians (people with maths degrees) is that they are a lot less strict about mathematical notation than you might expect. Their rule generally seems to be that clarity matters, while the exact syntax doesn't. This is something that I agree with.

For example, there are a certain set of people who will always write equations that involve division with horizontal bars, eg:

$$R_x = \frac{v_m R_1}{(V_S - v_m)}$$

This sort of notation definitely makes sense if you're writing an academic paper or working on a serious mathematical proof. However, that is for shorter formulae aimed at the average person. For that reason, in this magazine, we would be more likely to typeset that formula as $R_x = v_m R_1 / (V_s - v_m)$ or $R_x = v_m R_1 \div (V_s - v_m)$.

I prefer the latter rather than the former, simply because the operator ÷ is generally understood by everyone to mean only division (it had other meanings historically, but they are basically obsolete now). On the other hand, / can be a separator (either/or), especially in the context of computers, where it's often used as a directory separator, such as Unix-like systems.

I sometimes see arguments along the lines of "nobody uses ÷ to mean division after 5th grade", which I find odd for several reasons. One, the ÷ symbol for division is still used on a lot of calculators, both hardware and software types. Two, as I mentioned above, it's one of the clearest symbols we can use to indicate division apart from the horizontal bar, which takes a lot more space.

I think that most serious mathematicians would say that using ÷ is fine because of my first statement: that clarity matters more than syntax. Just to save a few pen strokes, most mathematicians won't use × for multiplication; they'll use a dot (·) or a space, or just write the two things to be multiplied next to each other!

They can get away with this for one simple reason: others still understand what they mean. For this reason, I don't think there's anything wrong with using any of these symbols, as long as it will be clear to the reader what you are using them for.

Mathematicians sometimes even redefine the meaning of certain operators for convenience! So I doubt they would turn their noses up at someone using ÷ in a simple formula if that's what the writer deems necessary.

One important aspect you may not have learned in school is that while the × and ÷ operators have the same priority, multiplication by juxtaposition has a higher priority. So in the case of $a \div bc$, you must multiply b and c before dividing the result into a.

Nicholas Vinen, Electron Publishing (Australia)*
* a division of Silicon Chip Publications Pty Ltd.

Practical Electronics

UK readers SAVE £1 on every issue

SUBSCRIBE NOW!

01202 087631

Take out a one-year subscription and save more than £10 over the year.
Even better – save over £1 per issue if you subscribe for two years – a total saving of £26.

Overseas rates represent exceptional value

You also:
- **Avoid any cover price increase for the duration of your subscription**
- **Get your magazine delivered to your door each month**
- **Guarantee a copy, even if your local newsagent sells out**

Order by:
- **Phone or post with a cheque, postal order or credit card**
- **Or order online at: www.electronpublishing.com**

Practical Electronics

SUBSCRIPTION PRICES

Subscriptions for delivery direct to any address in:

UK: 6-months £33.99, 12-months £59.99, 24-months £114.99

Europe Airmail: 6-months £38.99, 12-months £69.99, 24-months £129.99

Rest Of The World Airmail: 6-months £46.99, 12-months £84.99, 24-months £164.99

Cheques (in **£ sterling only**) payable to ***Select Publisher Services Ltd*** and sent to: Practical Electronics Subscriptions, PO Box 6337, Bournemouth BH1 9EH, United Kingdom

Tel: 01202 087631
Email: pesubs@selectps.com

Also via our online shop at**: www.electronpublishing.com**

Subscriptions start with the next available issue. We accept MasterCard or Visa.
(For past issues see the Back Issues page.)

DIGITAL SUBSCRIPTIONS

Electronic subscriptions are available from £29.99 for 12 months, more details at: **www.electronpublishing.com**

SUBSCRIPTION ORDER FORM

☐ **6 Months:** UK £33.99, Europe £38.99 (Airmail), Rest of the World £46.99 (Airmail)

☐ **1 Year:** UK £59.99, Europe £69.99 (Airmail), Rest of the World £84.99 (Airmail)

☐ **2 Years:** UK £114.99, Europe £129.99 (Airmail), Rest of the World £164.99 (Airmail)

To: Practical Electronics Subscriptions
PO Box 6337
Bournemouth BH1 9EH
United Kingdom

Tel: 01202 087631

Email: pesubs@selectps.com

I enclose payment of £ .....
(cheque/PO in £ sterling only)
payable to ***Select Publisher Services Ltd***

☐ Please charge my Visa/Mastercard

My card number is: .....
Please print clearly, and check that you have the number correct

Card Security Code ..... Valid From Date.....
(The last 3 digits on or just under the signature strip)

Card Ex. Date .....

Name .....

Address .....

Post code ..... Tel. .....

Email .....

Subscriptions start with the next available issue.

Buying Second-Hand SPEAKERS

Bargains can be had, but you need to know what to look for!

~ Feature by Julian Edgar ~

It's great fun building your own speakers, and excellent results can be gained from doing so. However, it is not a cheap process if you want good results. By the time you pay for new drivers, crossovers, wiring and connectors, then build the cabinets, the cost can really add up.

So what do you do if you're on a tight budget? The answer is to buy second-hand speakers. But how do you select speakers that are good, especially when you often cannot listen before buying?

Size

It's a fundamental fact of life that small loudspeakers will usually sound worse than large speakers. Yes, small enclosures can sound good – but it's technically much more difficult to achieve good outcomes, especially in the production of bass. Chances are then, the smaller the speaker, the worse it will sound.

Another reason for selecting larger speakers is efficiency – larger speakers are generally much more efficient than small speakers. What this means in practical terms is that a given sound volume will be achievable with less amplifier power. While the cost per watt of amplifier power has come down a lot in recent years, lower-power amplifiers still cost less than high-power amplifiers!

When browsing second-hand speakers, you'll find many home theatre systems that use five small speakers. Often, each about as big as a closed fist. Irrespective of their brand, don't buy these! The same applies to the tall, thin and shallow speakers often used as the front and rear speakers in home theatre systems. Again, avoid them.

About the smallest enclosure you can be confident will sound OK is what used to be called a large bookshelf speaker, say 20 litres in volume, using a 5-inch (12cm) or 6-inch (15cm) diameter woofer.

To roughly calculate the enclosure volume, multiply the height by width by depth in centimetres, divide by 1000 and subtract 20%. For example, if a cabinet is 40cm tall, 25cm wide and 30cm deep, the product of those is 30,000 (40 × 25 × 30). Dividing by 1000 gives you 30, then subtracting 25% (to account for the material thickness etc) gives you an estimated volume of 22.5 litres.

However, for main speakers, larger enclosures of around 40-60L will likely sound better.

Speakers are widely available second-hand – but how do you tell if they're any good, especially when often you cannot listen first? Tiny speakers invariably sound terrible, no matter the brand. With speakers, bigger usually equals better.

Buying Second-Hand Speakers

HAMMOND®

1591 ABS flame-retardant enclosures

Learn more:
hammondmfg.com/1591

uksales@hammfg.com • 01256 812812

Note that the shape of the enclosure doesn't matter nearly as much as the internal volume. If you need narrow-but-tall tower speakers to achieve your desired aesthetics, that's fine – as is an older, more squat enclosure.

Rule #1: bigger = better.

Weight

The next step in selection is to feel how heavy the speaker is for its size. Almost invariably, heavy speakers will sound better than lightweight designs. There are several reasons for this. Woofers with larger magnets are much heavier than those with small magnets; crossovers that use inductors and not just capacitors are heavier; and cabinets with thicker walls and internal bracings are heavier than those without.

So just by picking up the speakers, you will get a quick but usually quite reliable indicator of sound quality.

Rule #2: heavier = better

Enclosure design

Excluding rarer designs like folded horns and similar, enclosure design falls into two basic categories: sealed and ported.

A sealed design, as the name suggests, has no openings in the box – the air behind the drivers is trapped within the enclosure. The old name for this design approach – acoustic suspension – gives an indication of how the enclosure works. The springiness of the air within the enclosure provides the restoring force for the woofer's suspension.

When gently moved by your fingers, true acoustic suspension woofers tend to have a viscous, slow return movement. On the other hand, many cheaper sealed speakers use a conventional woofer that springs quickly back into position.

In general, unless it is a true acoustic suspension design, a sealed enclosure is likely to deliver worse bass than a ported enclosure. The exception to this is if the sealed enclosure is large compared to the size of the woofer – and the woofer is large as well!

A ported design has openings – ports – in the enclosure. These are usually on the front of the enclosure, but some speakers have them on the back. Always check both the front and the back panels for ports. Generally, a ported enclosure can achieve deeper bass from a given size of enclosure.

The rear of a lightweight speaker with a plastic enclosure – horrible! A pair of these speakers was priced quite high – many second-hand sellers have no idea of the value (or lack thereof) of the speakers they're selling.

Ports should be flared at both ends (not with square edges, as shown here) and have no internal steps. Ports also allow you to inspect the interior of the enclosure without removing a driver.

The woofer cone should be able to be gently pushed inwards without any binding or scratchy sounds.

Importantly for buyers of second-hand speakers, the port also allows you to see inside the enclosure without pulling out a driver. When assessing a ported speaker, use a torch and peer into the enclosure through the port. Normally, you'll be able to see either the inside of the back or front panels.

In both cases, if you can see bare chipboard, mark down the speaker. Instead, what you want to see is some type of acoustic absorbing material (foam or a fluffy material) covering the panel. This material stops internal sound reflections that tend to colour the midrange.

Another thing to look at when examining the port is its design and finish. A well-designed port has flares at both ends. This reduces port noise as the air moves back and forth in the port. Also, the inside of the port should be smooth – many ports on cheaper speakers have an external piece that connects to an internal cardboard tube, with a distinct step at the join.

Finally, be wary of speakers that have either very large, very small or very short ports – yes, all of these can be correctly used, but more often than not, they're giveaways of a poor design.

Rule #3: Prefer ported enclosures

Drivers

The next step is to examine the drivers carefully. Let's start with the woofer.

The woofer should have a long-travel suspension. Using the spread fingers of one hand, carefully push the cone inwards. The cone should move with no binding (a voice coil that catches often gives a scratchy feel and sound) and the edge suspension should not distort or collapse. Furthermore, the cone should not change in shape, ie, obviously flex.

The larger the woofer, the less distance it needs to be able to move. So if you're looking at an older speaker with a 12-inch (30cm) woofer, don't be concerned if it can be moved only 5mm – there will still be plenty of bass.

On the other hand, a 6-inch (15cm) woofer needs to have a lot of travel to be effective. The more woofers the enclosure has, the less travel each one needs.

Sharply tapping the ends of your fingers on the woofer cone will excite its resonant frequency – the resulting sound should be as deep as possible. This is an excellent test that evaluates both the woofer and its enclosure.

If the speaker is fitted with a large midrange speaker, again, very gently move its cone with your fingers. It will be able to be moved only a short distance, but again, there should be no binding.

The tweeter should preferably be a dome design – and don't try to move it! Some cone tweeters can sound fine, but on any speaker of the last 20 years, a cone tweeter usually indicates it is a cheaper, lower-quality design.

If you can, apply a 1.5V cell across the speaker terminals. The woofer should leap forwards or backwards (forwards, if you've observed the correct polarity). If you put your ear to the tweeter and midrange, you should also be able to hear sounds as the cell is connected and disconnected – especially if you draw the wire across the battery terminal to produce a scratchy noise.

This test shows that all the drivers are working.

Many second-hand speakers, especially at thrift shops and similar, have been damaged while on display. For example, the centre caps on woofers and dome tweeters can be pushed in. Woofer centre caps can usually be pulled out by using a fine needle or by careful use of a vacuum cleaner nozzle, but dome tweeters damaged in this way are usually ruined.

Don't buy speakers with non-removable grilles. The risk is simply too great; for example, the woofer's foam suspension may have completely perished.

Rule #4: carefully check the drivers

Speaker specifications

Many speakers have specifications written on the back of the enclosure. For brand name designs, you can do a quick Google search and find the same information. However, more often than not, this information is of little help – to put not a fine point on it, it's often garbage.

Really, only one specification is likely to be semi-accurate – and that's nominal impedance. If the plaque says "4 ohms", the AC impedance is likely to be near 4Ω.

Stated frequency response? It's often the complete stuff of fantasy – or the response was measured at ±12dB, making it equally useless! Power handling? Is it RMS or peak power? Or power on normal music material? Or power at great distortion just before the voice coil melts? Who knows?

Without the speaker efficiency being stated, the power figure gives you no idea of how loud the speaker can play anyway. Instead, when using it, you would simply turn it down when it starts to distort (or it's too loud for comfort).

Where they are available, look at the specifications, but don't expect to gain a lot from them. Personally, when buying second-hand speakers, I mainly just check the impedance, to ensure the designated amp can safely run them.

The centre dust cap in the left photo is damaged, although it could probably be pulled back into shape with the careful use of a vacuum cleaner nozzle. However, note how badly the outer suspension is distorting when the cone is pushed inwards. This is one to leave behind! Speakers like the one on the right with dome tweeters are preferable to those with cone designs.

Buying Second-Hand Speakers

Clues to quality

In addition to weight, there are other clues to quality. The first is the brand name.

This is a tricky area. Some well-known consumer electronics brands have speakers that range in quality from awful to very good, while some specialist speaker brands have reputations that are vastly overblown. Definitely don't buy on brand alone – always ensure that the above four rules are met by the speakers you're considering.

Another clue to quality is the terminal block – if there is one! Cheap and nasty speakers typically don't have a terminal block; instead, the cable just comes through a hole in the rear panel. Small spring clip terminals indicate the manufacturer has been cost-cutting – not very reassuring about what's inside the enclosure!

Better speakers use gold-plated binding posts and often can be bi-wired – that is, there are separate connections for the upper and lower frequencies, invariably bridged with plates. Any speaker with adjustable tweeter and/or midrange level controls is likely to be of better quality.

Editor's note: usually only relatively modern hifi speakers offer bi-wiring – plenty of excellent older speakers have a simple pair of binding posts. Bi-wiring is mostly just a fad anyway.

On some enclosures, you can see how the box was made. If that's the case, look for the thickness of the board and whether the tweeter has been recessed to give better phasing with the woofer – in fact, anything that shows that care was taken in the design and construction.

Be very wary of any speaker that looks home designed and built. It may be a brilliant design – but more often, it's been built by someone using rules of thumb... or no acoustic design at all (if they used quality drivers, such speakers can make good donors for a new design).

Do not buy speakers where some or all the drivers have obviously been replaced. Many people buy replacement drivers based only on size, not the acoustic match for the enclosure and/or crossovers.

Speakers that use only one full-range driver are typically quite inferior. Many are at the cheap and nasty end of the range, and those from higher-quality brands usually require dedicated amplifier equalisation to sound even half reasonable.

Rule #5: look for quality clues

Passive radiators

There is an enclosure design that, in effect, combines both sealed and ported approaches. This design uses a passive radiator, which is like a second woofer, but it does not have a magnet or coil. The passive radiator moves back and forth like the air in a port but without ever fully unloading the woofer.

Some passive radiators are obvious (eg, they have a flat panel), while others look just like a second woofer. Passive radiator designs have the theoretical ability to develop excellent bass from small enclosures, but unfortunately no second-hand passive radiator speakers I've ever bought have sounded very good!

Conclusion

I am a bit of a nut when it comes to speakers – I have bought literally dozens of pairs over the last 50 years. What I have found is that if the speaker is large, heavy and ported, the drivers are in good condition, and the enclosure design and construction look fine, the chances are very much that you have a quality pair of speakers. PE

ENIGMA
COMPONENT SHOP
www.enigma-shop.com
Electronic Components for the Hobbyist!
Cable & Wire • Capacitors • Connectors • Filters
Kits & Modules • Hardware & Cooling
and much more...
READY TO BE REWARDED?
Join our loyalty scheme and earn points on every purchase you make which can be used for discounts in the future.
Read our reviews on Trustpilot

By Andrew Levido

Power Electronics

Part 7: Resonant Converters & Soft Switching

Higher switching frequencies can make input and output filters simpler, with smaller magnetics. They also allow a faster response to changes in the load current. Switching losses become a major problem at higher frequencies, but there is a solution.

We saw last month that switching losses in power electronic converters can become dominant as switching frequencies increase. However, higher frequencies are desirable as they allow the designer to increase the bandwidth of the control system, so it can respond to load changes more quickly.

We will start with a quick review of how a Mosfet switches on and off. The upper-left part of Fig.1 shows a typical boost converter. We will assume this is operating in periodic steady-state and that the inductor current is more-or-less constant at the timescale of the switch-on and switch-off of the Mosfet, typically in the 10ns range.

We will also assume that the Mosfet is off and the full load current is flowing through the diode at the instant of switch-off.

I have also drawn the gate equivalent circuit of the Mosfet below the boost converter circuit. When the device is off, the gate voltage is zero, the gate-source capacitance (C_{GS}) is fully discharged and the gate-drain capacitance (C_{GD}) is charged to the drain voltage, v_D. The gate-drain capacitance is much smaller than the gate-source capacitance, but it plays a big part in switching losses, as we shall see.

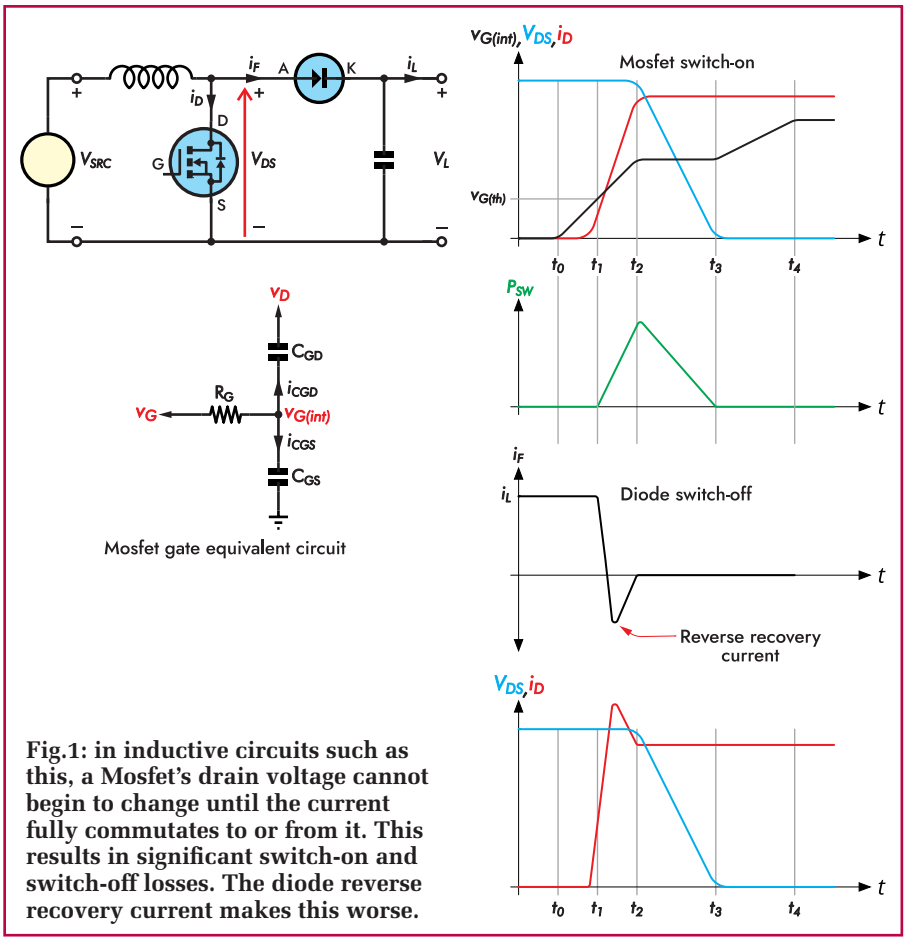

Fig.1: in inductive circuits such as this, a Mosfet's drain voltage cannot begin to change until the current fully commutates to or from it. This results in significant switch-on and switch-off losses. The diode reverse recovery current makes this worse.

The resistance R_G is the combination of the internal gate resistance and the source impedance of the Mosfet driver. The voltage $v_{G(INT)}$ represents the voltage at the gate metallisation on the Mosfet die that modulates the conductivity of the channel.

The plot labelled "Mosfet switch-on" shows what happens when we switch the Mosfet on. When the gate voltage is applied at time t_0, nothing happens immediately because the internal gate voltage (black trace) has yet to charge to the switch-on threshold.

Once the threshold voltage is reached at time t_1, the drain current (red trace) begins to rise. Because the load is inductive, the drain current must rise to its full extent before the drain-source voltage can begin to fall, at time t_2.

You can understand why this is the case with reference to the boost converter schematic. When the Mosfet starts to conduct, the current shifts (commutates) from the diode to the Mosfet. Until the Mosfet takes over 100% of the inductor current, the balance continues to flow through the diode, keeping the Mosfet's drain voltage fixed at the converter's output voltage.

This same phenomenon occurs in many converter types, including buck converters, but it's easiest to visualise with the boost converter since the Mosfet is grounded.

The rate at which the drain voltage can fall is determined by how fast C_{GD} can be discharged. The only thing discharging C_{GD} is the current iC_{GD} provided by the gate drive. While v_{DS} is falling, all the gate current is diverted into C_{GD} due to the Miller effect, so the internal gate voltage remains essentially constant until the drain-source voltage reaches (almost) zero at time t_3.

This 'flat spot' on the internal gate voltage is known as the Miller plateau.

After t_4, the two gate capacitances are effectively in parallel, and the

gate voltage continues to rise to its final value at t_4, where the channel is fully enhanced and the Mosfet's on-resistance is minimised. Switch-off is basically the reverse of this process.

The importance of all this is that there is a period, from t_1 to t_3, where there is significant voltage across and current flowing through the device at the same time, and therefore considerable power dissipated during the relatively short switch-on and switch-off periods. In fact, things are worse than I just described if we take into account the 'switch-off' characteristics of the diode.

When a diode switches from the conduction state to the blocking state, it does not switch off instantaneously. A large reverse current flows for a short period while the diode recovers its blocking capability – shown in the "Diode switch-off" plot.

This 'reverse recovery' current occurs because the majority carriers stored in the PN junction have to be extracted when the diode is reverse-biased. The amount of this charge (Q_{rr} in the data sheets) is small, but because it moves very quickly, the peak current can be high. This does not have a huge impact on the diode losses, but can contribute significantly to Mosfet losses.

When the Mosfet switches on at time t_2 and the inductor current commutates from the diode, the Mosfet sees an additional current spike due to the diode's reverse recovery (bottom chart). This occurs while the drain-source voltage is still high, so it adds to the Mosfet switch-on losses.

While it is often convenient to think of power Mosfets as voltage-driven devices, the description above demonstrates the importance of gate current in the switching process. The rate of change of drain-source voltage (dv/dt) during switching depends on the gate-drain capacitance and the gate current that charges and discharges it.

You generally need to drive the gate hard if you want to increase the dv/dt and minimise switching loss. However, a higher dv/dt produces in significantly more conducted and radiated EMI, so finding a compromise is usually necessary.

How significant are these switching losses? I made a simulation of the boost converter circuit using the SiHA120N60E Mosfet. This is a 650V, 25A-rated device in a TO-220 package. I set the boost converter input voltage to 200V, the output voltage to 400V and the load current to a slightly unrealistic 20A. I drove the gate to 15V via a 10Ω gate resistor.

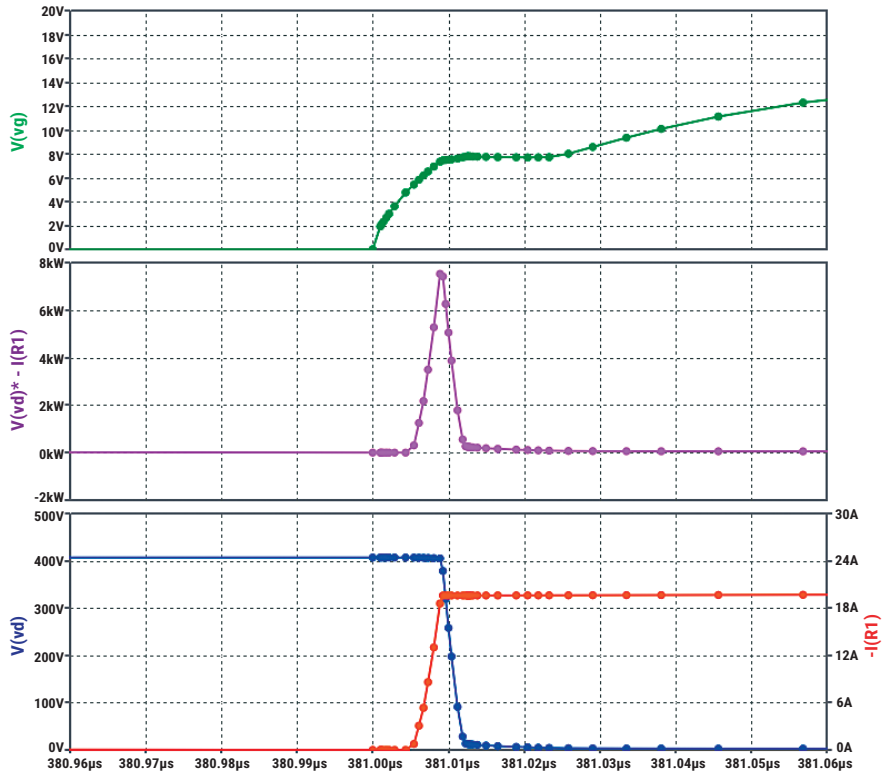

Fig.2: this simulation of the circuit in Fig.1 uses the SiHA120N60E Mosfet switching 400V at 20A. The switch-on losses peak at 7.5kW, although only for a few nanoseconds.

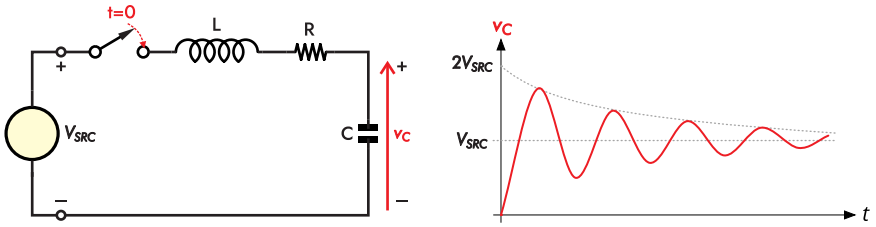

Fig.3: a resonant circuit like this can be described by two quantities: the natural frequency and the damping factor. We often use the 'quality factor' or Q to describe the relationship between the two.

The simulated switching waveforms are shown in Fig.2. You can clearly see that the drain current (red) rises fully before the drain-source voltage (blue) can begin to fall. You can also see the Miller plateau in the gate voltage (green). The purple trace is the instantaneous power dissipation in the Mosfet. It peaks at about 7.5kW, although the whole spike is only about eight nanoseconds long.

The total energy dissipated at switch-on is about 30µJ. If we assume the same for switch off, the total will be 60µJ per cycle. At 10kHz, this corresponds to a modest 600mW in losses, but if we want to switch at 1MHz, we are looking at switching losses of 60W – a much less attractive proposition!

Resonant circuits

Since switching losses are the product of voltage across and current through the switch, one way to reduce or eliminate switching losses would be to ensure that one or both of these quantities is zero at the time of switching. This type of switching is sometimes called zero-voltage switching (ZVS), zero-current switching (ZCS) or just described by the generic term 'soft switching'.

The usual way to ensure that voltage or current is zero when we switch is to exploit resonance. For a quick refresher on resonance, take a look at Fig.3, which shows a simple RLC filter with

a DC source and a switch that closes at time zero.

When the switch initially closes, a current will build up in the inductor, charging the capacitor through the resistor. The voltage on the capacitor will continue to rise past V_{SRC} to the point where the current in the inductor reverses and it begins to fall. When the capacitor voltage falls low enough that the inductor current reverses again, it starts to rise.

This oscillation continues, but is damped by the resistance. Eventually, the capacitor voltage settles at V_{SRC} and the inductor current goes to zero.

This is a damped oscillation. If the resistance were zero, the oscillations would continue indefinitely (in theory). If the resistor had a very high value, the circuit would behave like a standard RC filter, with the capacitor voltage rising smoothly (and exponentially) to V_{SRC}.

We can therefore describe these resonant circuits with two quantities: the natural frequency and the damping factor. The natural frequency, designated ω_0, is the frequency at which the undamped LC network would oscillate. This is given by $\omega_0 = 1 \div \sqrt{LC}$.

The damping factor, designated by the Greek letter α, is equal to $R \div 2L$. The damped circuit will oscillate at a frequency lower than the natural frequency. This frequency, ω_d, called the natural frequency, is equal to $\sqrt{\omega_0^2 - \alpha^2}$. These frequencies are expressed in radians per second, where 2π radians per second equals 1Hz.

We don't generally use the damping factor directly in our design process. Instead, we use a related quantity, the 'quality factor' or Q of the circuit. Q relates the damping factor to the natural frequency by the expression $Q = \omega_0 \div 2\alpha$. A resonant circuit with high Q has low damping, and vice versa.

You can play around with substituting these expressions into each other and get two other useful definitions: $Q = \omega_0 L \div R$ and $Q = 1 \div \omega_0 CR$.

We can use this information to build a resonant DC-AC converter, as shown in Fig.4. A DC source feeds a half-bridge switch followed by an LC series filter and a resistive load. The end of the load is held at $\frac{1}{2}V_{SRC}$ by the two bypass capacitors. These have a value large enough that their midpoint voltage remains more-or-less constant at the switching frequency.

The voltage across the filter and load is therefore a square wave with amplitude $\frac{1}{2}V_{SRC}$, as shown in the red trace. If we switch the converter at a frequency equal to the damped natural frequency of the filter and load, the load current and voltage will be a relatively pure sinusoid at that frequency.

With the switching frequency equal to the natural (resonant) frequency, the inductive and capacitive reactances cancel out, so the load looks resistive and the maximum possible energy is transferred to it. We say that the voltage gain of the resonant tank is unity under these conditions.

By this, we mean that the AC voltage across the resistive load (blue) is equal to the amplitude of the fundamental of the square wave. This gain is called the 'tank gain'.

If the switching frequency is a little higher or lower than the natural frequency, the circuit looks a bit inductive or capacitive, respectively. The tank gain falls off either way, and the amount of power transferred to the load reduces. We can therefore control the load power by frequency-modulating the drive signal.

We can choose to use either 'above resonance' or 'below resonance' control strategies. We will see how this works in practice later.

The graph in Fig.4 suggests that switching occurs at the current zero-crossing, but this is a bit misleading. It is true that the current is near-zero at the time of switching if the switching frequency is precisely aligned with the resonant frequency, but we have already discussed that we will operate at a higher or lower frequency to control the output power.

However, it is easy enough to modify this circuit to achieve zero-voltage switching. It just requires the addition of a small capacitor across each switch and a short 'dead time' during which both switches are off. This is shown in Fig.5. Here, we are using over-resonance control, so the LC filter looks inductive and the current lags the voltage by the angle θ.

The charts to the right of the figure show the load current (blue) and the voltage across the lower switch, S2 (red). During period A, the upper switch S1 is conducting, so C1 is discharged and C2 is charged to V_{SRC}.

At the beginning of period B, S1 opens while the load current is still positive. Capacitors C1 and C2 take over providing the load current, and the voltage across S2 falls while C1 charges and C2 discharges.

If the capacitor values are chosen such that the period B is large with respect to the (~10 nanosecond) switching time, the voltage across S1 at the

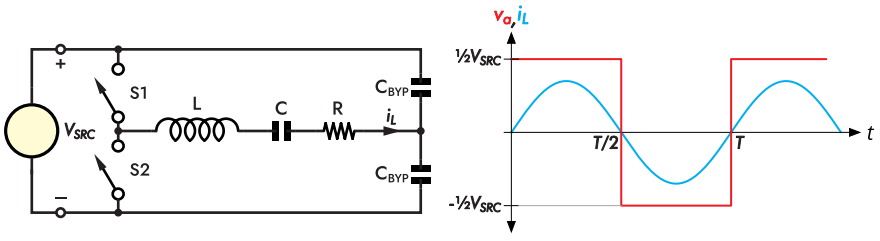

Fig.4: if we switch this resonant DC-AC converter at its natural frequency, the inductive and capacitive reactances cancel, and maximum power is transferred to the resistive load. By shifting the frequency above or below the natural frequency, we can control the power to the load.

5-year collections
2019-2023
£49.95
2018-2022
£49.95
2017-2021
£49.95
2016-2020
£44.95
Purchase and download at:
www.electronpublishing.com

time it opens will be effectively zero – C1 will hold the voltage across S1 to zero while it opens.

At the start of period C, C1 is fully charged to V_{SRC} and C2 is fully discharged. The still-positive load current now commutates to the freewheeling diode, D2. At this time, S2 can be closed while the voltage across it is zero. When the load current reverses at the start of period D, S2 is closed, ready to take the load current for the bulk of its negative excursion.

At the start of period E, S2 is opened, while C2 holds the voltage across it to zero. The capacitors support the load current until the start of period F, when the freewheeling diode, D1, takes over. This is when S1 is closed, while its drain-source voltage is zero.

The upshot of all this is that both switches only ever open or close with zero voltage across them, resulting in very low switching loss. The caveats to this are that the device switch-off time is small compared to the charge/discharge times of C1 and C2, and that there is enough phase lag that the freewheel diodes are conducting when the switches are on.

The former is not such a challenge, since the switching time is short, but the latter means that we cannot operate too close to natural frequency so the filter inductance remains high enough.

Resonant DC-AC converters like this are pretty common – for example, most induction cooktops work this way. A typical large domestic induction cooktop contains resonant converters capable of a power output up to 7kW at frequencies in the 20-100kHz range. This would not be practical without zero-voltage switching.

Resonant DC-DC converters

You could also imagine rectifying and filtering the output of a resonant DC-AC converter, perhaps after passing it through a transformer, to produce a DC output, as in Fig.6.

Frequency modulation could be used to control the resulting DC output voltage.

We can think of this converter as a series of four blocks, each with its own voltage gain. The product of these gains is the voltage transfer function of the converter: $V_L \div V_{SRC} = G_i G_t G_x G_r$.

Calculating some of these gains is easy. The inverter puts out a square

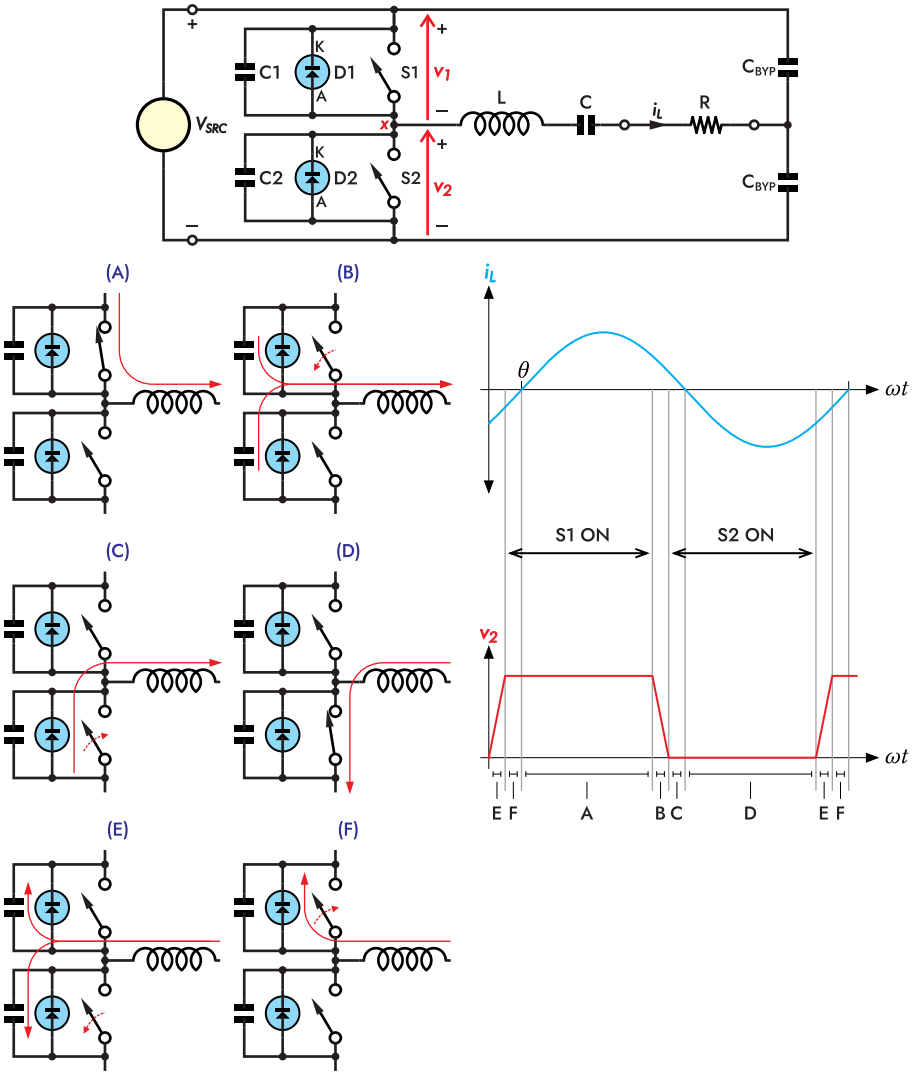

Fig.5: with the addition of capacitors across the switches and a short dead-time when both switches are off, we can achieve almost lossless switching. The switching frequency must always be a little above the natural frequency for this to work.

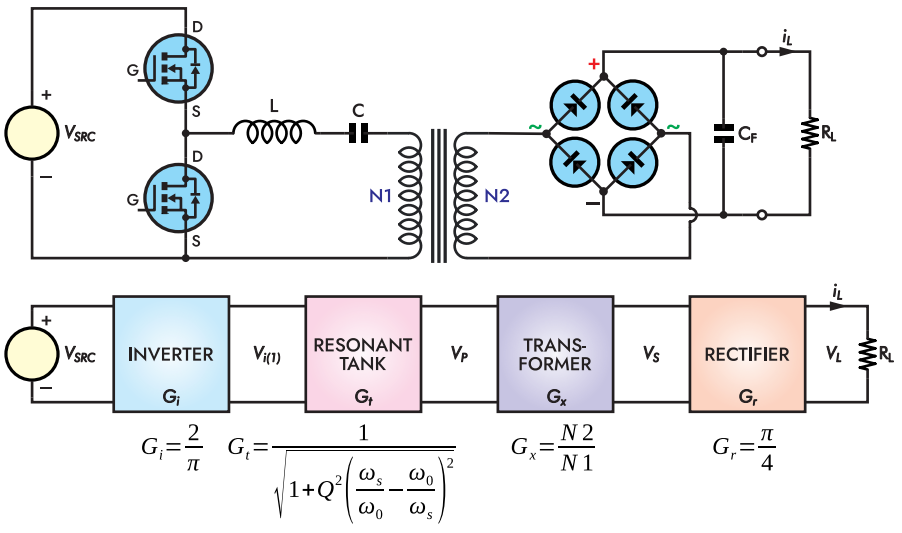

Fig.6: a resonant DC-DC converter can be thought of as four distinct blocks, each with its own gain. This simplifies the analysis enormously.

wave with a peak-to-peak amplitude of V_{SRC}, so an amplitude of $\frac{1}{2}V_{SRC}$. The current is sinusoidal, so only the fundamental component of this voltage can transfer real power.

We have seen many times before that the amplitude of the fundamental frequency of a square wave is $4 \div \pi$ times its amplitude. Since the amplitude is $\frac{1}{2}V_{SRC}$, the gain through the inverter must be $G_i = 2 \div \pi$.

The gain of the tank is much more complex. I won't go through the derivation (life is too short as it is), but it can be shown to be the ugly expression under the resonant tank block in Fig.6. The important thing to note is that the tank gain depends on the ratio of the damped-to-undamped natural frequencies and the Q of the tank. We'll look into this more when we get to an example.

The transformer's voltage gain is trivial to calculate – it is just the turns ratio, as you would expect. Calculating the rectifier's gain is a bit harder, but not much. Fig.7 shows how.

The resonant tank current driving the transformer primary will be sinusoidal, so we can model the transformer's secondary current as a sinusoidal current source, i_S, with some amplitude I, which will depend on the primary current and the turns ratio. The waveforms associated with this simplified circuit are shown on its right.

If we assume the filter capacitor is large enough to make the voltage ripple negligible and the diodes are ideal, the transformer secondary voltage v_S will be a square wave with amplitude V_L. Because i_S is sinusoidal, only the fundamental component of v_S can contribute real power to the load.

Again using the relationship for the fundamental of a square wave, the gain of the circuit $V_L/v_{S(1)}$ is $\pi \div 4$.

We also need to work out the equivalent AC resistance of the rectifier and load. This is important because it is this resistance, seen through the transformer, that loads the resonant tank and determines its damped natural frequency and Q, both of which will affect the tank gain.

Dividing $v_{S(1)}$ by the secondary current gives an equivalent AC resistance of the rectifier and filter of $R_{(ac)} = (4 \div \pi) \times (V_L \div I)$. Expressed in terms of the converter's power output, $R_{(ac)} = (8 \div \pi^2) \times (V_L^2 \div P)$. Noting that the last bracketed term is equal to the load resistance, $R_{(ac)} = (8 \div \pi^2) \times R_L$.

A practical example

We now have all the equations we need to look at a practical design example. There is not enough space here for a comprehensive design exercise, but I want to show how one would approach such a design.

Let's imagine we are building an isolated resonant DC-DC converter to operate from rectified mains and deliver 10V DC at 20A (so 200W) into a resistive load. The input voltage range should be 300-400V to accommodate a range of mains voltages (let's not worry about supporting 110-120V AC mains just yet).

Because we are using the frequency to control the output power, we need to specify a minimum load so the frequency range is bounded at both ends. We will use a minimum load of 5A (50W) for this exercise. We will use above-resonance control with a target switching frequency in the range of 500kHz to 1.5MHz, or thereabouts. I will break the design up into steps.

1. We can start by calculating the maximum and minimum load resistances corresponding to the minimum and maximum output currents: $R_{L(MIN)} = 0.5\Omega$ and $R_{L(MAX)} = 2\Omega$. We can also choose an undamped natural frequency at the low end of our desired range, say 600kHz or 3.77×10^6 radians per second.

2. We have to design the transformer turns ratio so we can achieve the desired output voltage when the source voltage is at its minimum. The minimum DC voltage times the inverter gain G_i gives us a minimum input voltage of $191.0V_{RMS}$. On the other side of the transformer, the 10V output voltage divided by the rectifier gain G_r tells us that the fundamental of the secondary voltage must be $12.7V_{RMS}$.

The ratio of these values gives us a transformer gain G_x (secondary to primary turns ratio) of 0.067. We actually need a bit more gain than this because I have neglected the rectifier diode drops and any other losses. I

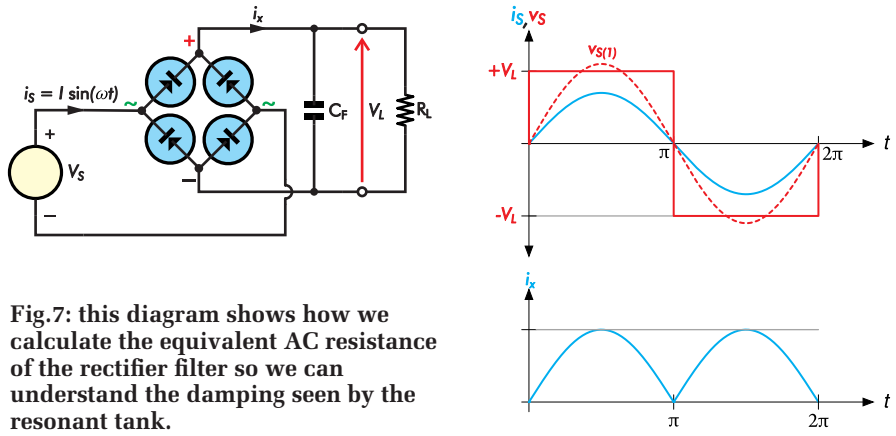

Fig.7: this diagram shows how we calculate the equivalent AC resistance of the rectifier filter so we can understand the damping seen by the resonant tank.

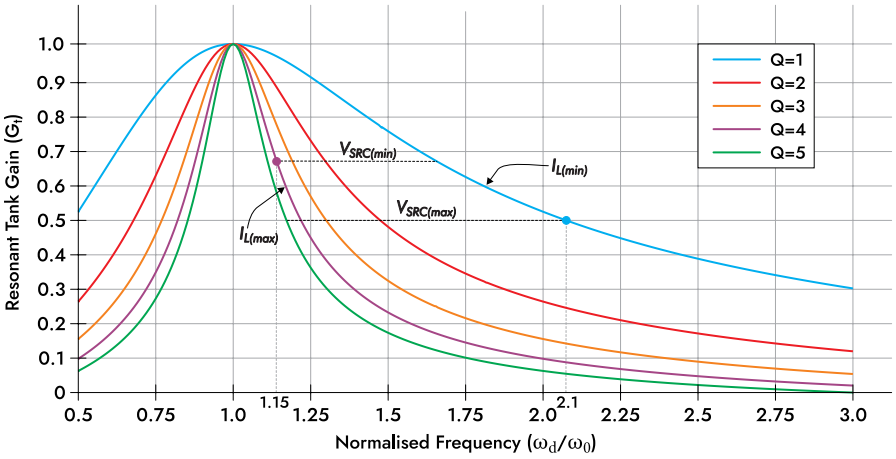

Fig.8: the curves show the tank gain vs normalised frequency for various values of Q. The example in the text operates in the region bounded by the dotted lines and the Q=1 and Q=4 curves.

will therefore use a nice round turns ratio of 0.1 (ten primary turns for every secondary turn).

3. Now we can calculate the required tank gain. This is easy because we know the inverter gain, the transformer gain, the rectifier gain and the required end-to-end voltage gain (V_L/V_{SRC}). Because we have a range of input voltages, we will also have a range of tank gains. It turns out that the tank gain G_t has to range between 0.5 and 0.67.

4. The last thing we have to do before we can calculate the component values is to work out what the load resistance looks like from the primary side of the transformer. This is the resistance that will load the resonant circuit.

We saw from the analysis of the rectifier that the secondary-side AC resistance of the load is $(8 \div \pi^2) \times R_L$. This transforms our 0.5Ω and 2Ω minimum and maximum load resistances to 0.405Ω and 1.62Ω respectively.

We then have to reflect these resistances through the transformer ratio by multiplying them by $(N_1 \div N_2)^2$, which just means multiplying them by 100 in our case. The effective resistance loading the tank is therefore in the range of 40.5Ω to 162Ω.

5. We calculate the resonant tank component values based on the Q. The minimum Q occurs when damping is highest and the resistance is at its maximum, corresponding to light loading on the converter. We can just choose the minimum Q to be 1 and calculate the resonant inductor from $Q = (\omega_0 L) \div R$.

Rearranging to make L the subject and plugging in the other values ($\omega_0 = 3.77 \times 10^6$ radians per second and R = 162Ω) gives an inductance of 43.0μH. We can then calculate C from the relationship $\omega_0 = 1 \div \sqrt{LC}$ to give 1.63nF.

6. Finally, we can calculate the maximum Q, which occurs when the load is heaviest and the resistance load on the tank is lowest. We can use the same $Q = (\omega_0 L) \div R$ formula, this time plugging in the inductance we just calculated and the 40.5Ω minimum resistance. This gives us a maximum Q of 4, which is not unreasonably high.

7. You could use the ugly formula for tank gain in Fig.6 to calculate what this means for the damped natural frequency, but it is probably easier to follow if you look at the graph in Fig.8. This plots the tank gain vs the normalised frequency (the ratio of damped natural frequency to the natural frequency) for various values of Q.

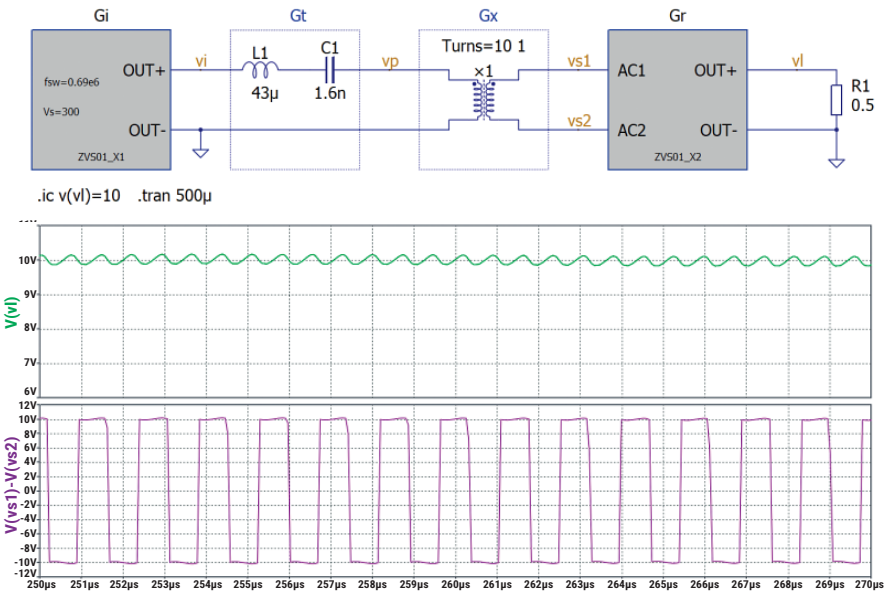

Fig.9: this simulation of the example resonant DC-DC converter agrees with the calculations. The upper green trace is the output voltage and the lower mauve one is the transformer secondary voltage.

Our converter will operate in the region bounded by the two horizontal dotted lines (tank gain of 0.5 to 0.67) and the curves corresponding to Q=4 (purple) and Q=1 (blue/cyan). The tank gains correspond to the range of input voltage and the Q values correspond to the load resistance range.

We can then read off the minimum and maximum normalised frequencies from the horizontal axis. I have marked these points with large dots. In this example, we expect the resonant frequency to range from $1.15\omega_0$ to $2.1\omega_0$. This corresponds to a frequency range of 690kHz to 1.24MHz.

Results

I find this type of graph very intuitive. The lowest switching frequency corresponds to a heavy load, low input voltage scenario. The highest switching frequency corresponds to the lightest load and highest input voltage scenario. You can see now why we have to set a minimum load – the switching frequency will go through the stratosphere if the load resistance gets too high.

I could not resist simulating this circuit, as shown in Fig.9. The inverter block on the left just contains a behavioural voltage source that produces a 50% duty cycle square wave with the frequency and amplitude specified on the front. The resonant tank and transformer are obvious, and the rectifier block consists of a full-bridge of ideal diodes and a 10μF capacitor.

The simulation run below the schematic was at the highest load, lowest input voltage operating point. The results are a bit unspectacular, with the DC output voltage in green and the transformer secondary voltage in purple. However, it does confirm that this switching frequency is roughly correct to achieve the output voltage we desire.

There is obviously a lot more to designing a resonant converter, especially one at this power level. In fact, this article has just given a small introduction to resonant and soft switching converters; there are countless variations out there, including some quite novel and interesting circuits.

Conclusion

This article concludes our series on power electronics. We have covered a lot of ground, including DC-DC, AC-DC and DC-AC converters. We have touched on control systems, magnetics and EMI filtering, and with this article, resonant converters.

As I stated at the outset, this series was not meant to be a university-style course on power electronics. Rather, I hope I have provided some insights and a few tools and techniques that may be useful in exploring this endlessly fascinating topic. PE

By Tim Blythman

Decoder / Base Station / Using DCC / Remote Controller / DCC Booster

DCC Decoder

for model locomotives

DCC is a great way to control model trains in a realistic fashion. You need a DCC Decoder for every locomotive and a base station to produce DCC signals on the track. This DCC Decoder project provides the first essential part. We will follow up with a DCC Base Station project and a guide to getting started with DCC using these in our next issue.

Image source: https://pixabay.com/photos/steam-locomotive-model-series-23-3627896/

For many years, model railways were limited to a simple DC system. The two rails formed conductors that carried power from a controller directly to the motor of a model railway locomotive. The controller would simply provide a voltage and polarity to control the loco's speed and direction. It's simple and elegant, but only allows one train to be operated at a time.

Of course, real railways run more than one train simultaneously and, over time, different technologies have been used to allow this. One well-known example divides the track into sections called blocks. A switch panel routes power between the blocks and multiple controllers. As you can imagine, this can get very complex.

In the late 1980s, Lenz Elektronik of Germany developed the early versions of what would later become digital command control or DCC. The US National Model Railroad Association (NMRA) took an interest and developed DCC as a standard under its auspices.

The NMRA had previously worked on standards relating to matters such as wheels and trackwork for model railways, intending to foster compatibility and interoperability between manufacturers. With its promotion as a standard, DCC continues to be one of the most popular systems for operating model railways.

The NMRA continues to develop and publish updates to the DCC standards and these, along with their other standards are available to download from https://pemag.au/link/ac7w

We have published several articles related to DCC, including feature articles and projects for boosters, programmers and even a computer-based base station. However, this is the first DCC decoder we have published, which means that all you need to build a complete DCC system can now be found within *Practical Electronics*.

The Complete Arduino DCC Controller can round out a working system (as described in the January 2021 issue), but it relies on a computer, so might be excessive for readers looking for a straightforward starter system. Thus, we have also designed a standalone Base Station based on a Raspberry Pi Pico 2 microcontroller and touch panel LCD, which we will introduce next month.

That unit can control up to five locomotives when paired with a suitable power supply, and can also perform decoder programming; more than enough for most small layouts. If you are not familiar with DCC terminology, refer to our glossary later in this article.

DCC operation

Fig.1 is a comparison between so-called DC operation (using a controller providing an analog DC voltage) and DCC operation. Fig.1(a) shows DC operation.

This project deals with the decoder shown in Fig.1(b). Its main role is receiving power and data as a varying voltage sent from the DCC base station, and decoding the data to drive the motor and function outputs. The latter might be incandescent globes or LEDs. Note that the motor and lights are not part of this project; they will be part of the locomotive to which you are adding the Decoder!

Our Decoder and Base Station can

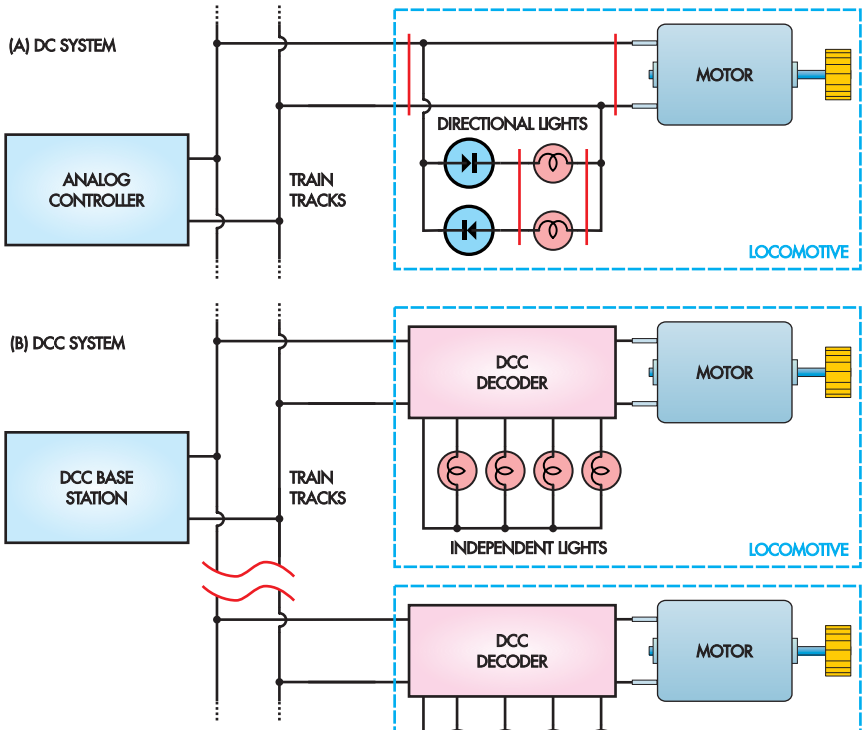

Fig.1: (a) shows the typical wiring inside a DC-only (non-DCC) locomotive, while (b) shows how this changes when a DCC decoder is fitted. If two locomotives were on the track in (a), their motors would run together. The decoder allows the lights and motors in the two locomotives to run independently. The vertical red lines show where wires need to be cut to install the Decoder.

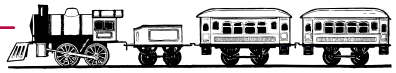

Features & Specifications

- Suitable for N and HO/OO scales
- Has support for a keep-alive capacitor
- Long, short and consist addressing
- Acceleration control and speed curve adjustment
- Configurable lighting outputs and flashing effects
- Voltage compensation option
- Supports operations, paged, physical and direct programming modes
- Supports 14-step, 28-step and 128-step speed control
- Robust design with current limiting on all outputs
- Size: 28 × 14 × 4mm
- Input voltage: up to 17V peak
- Motor current: up to 500mA
- Function outputs: four, up to 100mA each

be the electronic basis of an economical but complete DCC system. The Decoder design is robust, with current limiting provided for both the motor and function outputs.

Commercial DCC decoders include a multitude of features, including sampled and synthesised sounds and complex lighting effects. Our Decoder provides some basic lighting effects and is an inexpensive option if you assemble it yourself.

We've also designed the Decoder with the idea that our readers might want to extend or modify it by adding circuitry. Thus, the Decoder provides connections for rectified DC from the track as well as regulated 3.3V, something not commonly seen on commercial decoders.

Alternatives

These days, RF (radio frequency) controls have become more accessible and microcontrollers more powerful. We have published a few projects relating to RF control of model railways, including Les Kerr's Battery-Powered Model Train from the December 2025 issue.

We have also seen systems based on WiFi. These typically use low-cost ESP8266 or ESP32 WiFi processor modules. Still, we are yet to see uniform standards for these sorts of systems, so you will likely be committed to building all the parts yourself if you choose one of these options.

Our DCC Decoder and Base Station have both been designed according to the DCC standards, and tested with their commercial counterparts. So they can be used as part of an existing DCC system, or further expanded in the future with other products.

The Decoder

A locomotive decoder (also known as a mobile decoder) is intended to be used with a model railway locomotive or perhaps a motorised railcar. It could be fitted into the tender behind a steam locomotive, if there is space, or if that is where the motor is installed.

There are other types of decoders, such as accessory decoders, that can operate line-side equipment, such as points motors and signal lights. An accessory decoder could even control lights or other effects in model buildings and such.

A function-only decoder is another type of decoder that does not drive a motor, but can operate lights or other effects. A typical use would be for internal or marker lights in a carriage, or perhaps to control a smoke generator in a guard's van to emulate a wood stove.

Since a locomotive decoder can also control these functions, there is no reason our Decoder cannot be used as

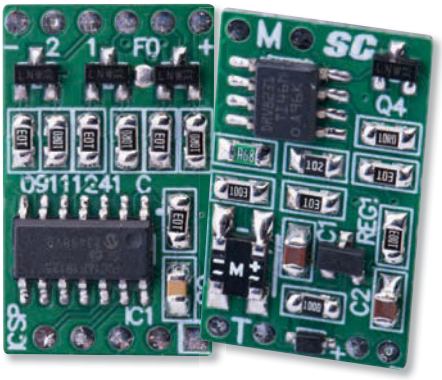

This tiny DCC Decoder board can be fitted to a model railway locomotive, allowing it to be controlled by a DCC Base station. It's only 14mm wide (shown here at twice actual size), so is suitable for HO/OO and N scale models.

a function-only decoder in other rolling stock. Still, we expect that many readers new to DCC will be most interested in adding the Decoder to their locomotives as a starting point.

As a rough guide to getting started, you will need a locomotive decoder for each locomotive or other item of motorised rolling stock, plus a base station to coordinate them all. This article will present the Decoder design; for those who already operate a DCC system, we'll provide a wiring guide and some information about its operating parameters; particularly its configuration variables (CVs).

This information should be enough to allow our Decoders to be customised and integrated into an existing DCC layout. For those new to DCC, we'll provide a more complete user guide alongside the DCC Base Station in the next article.

Circuit

Fig.2 shows the circuit of our four-function DCC mobile Decoder. Some of the connections are marked with colours; these correspond to the wire colours used by the DCC standards. Each connector on the Decoder is just a PCB solder pad. With space at a premium, the expectation is that wires are soldered directly to the Decoder PCB.

The DCC track signal is an AC square wave of about 12V-15V in amplitude (24V-30V peak-to-peak) with its frequency varying around 6kHz. The actual voltage is not too important, although our design limits it to about 7V minimum and 17V maximum. Just over 12V is typical, since this is the nominal operating voltage of the motors in many locomotives.

The track connections to the Decoder (red and black) feed into bridge rectifier BR1 and then REG1, an MCP1703 3.3V LDO (low-dropout) regulator, providing our nominal 12V and 3.3V rails. The MCP1703 has a maximum operating input voltage of 16V; our 17V maximum track voltage is due to this, with a drop of at least 1V due to the bridge rectifier.

Note that some DCC systems operate at up to 22V peak, so you should check this before using our Decoder with an existing system. Some systems can provide a wide range of supply voltages, so it may be possible to simply reduce it.

The 3.3V rail powers IC1, a PIC16F18126 microcontroller, at its pins 1 and 14. This is the 14-pin sibling of the 20-pin PIC16F18146 micro we have used in numerous other projects. The firmware we have created will work in a PIC16F18124 or PIC16F18125 (with less memory), but we will supply the PIC16F18126 parts in the kit as it's only slightly more expensive.

The three capacitors provide the necessary bypassing for the regulator and microcontroller, while a pair of 100kΩ resistors connect pins 10 and 11 of the micro to the track connections. The 100kΩ resistors limit current flowing into the pins, and allow the micro to safely detect the polarity of the higher track voltage and thus receive the DCC signals.

An 11:1 divider (10kΩ/1kΩ) on the 12V rail connects to pin 9 on IC1, allowing the voltage on this rail to be measured by the 3.3V micro.

IC2 is a DRV8231 3A full-bridge motor driver IC, powered at its pins 1 and 5 from the 12V rail. Digital

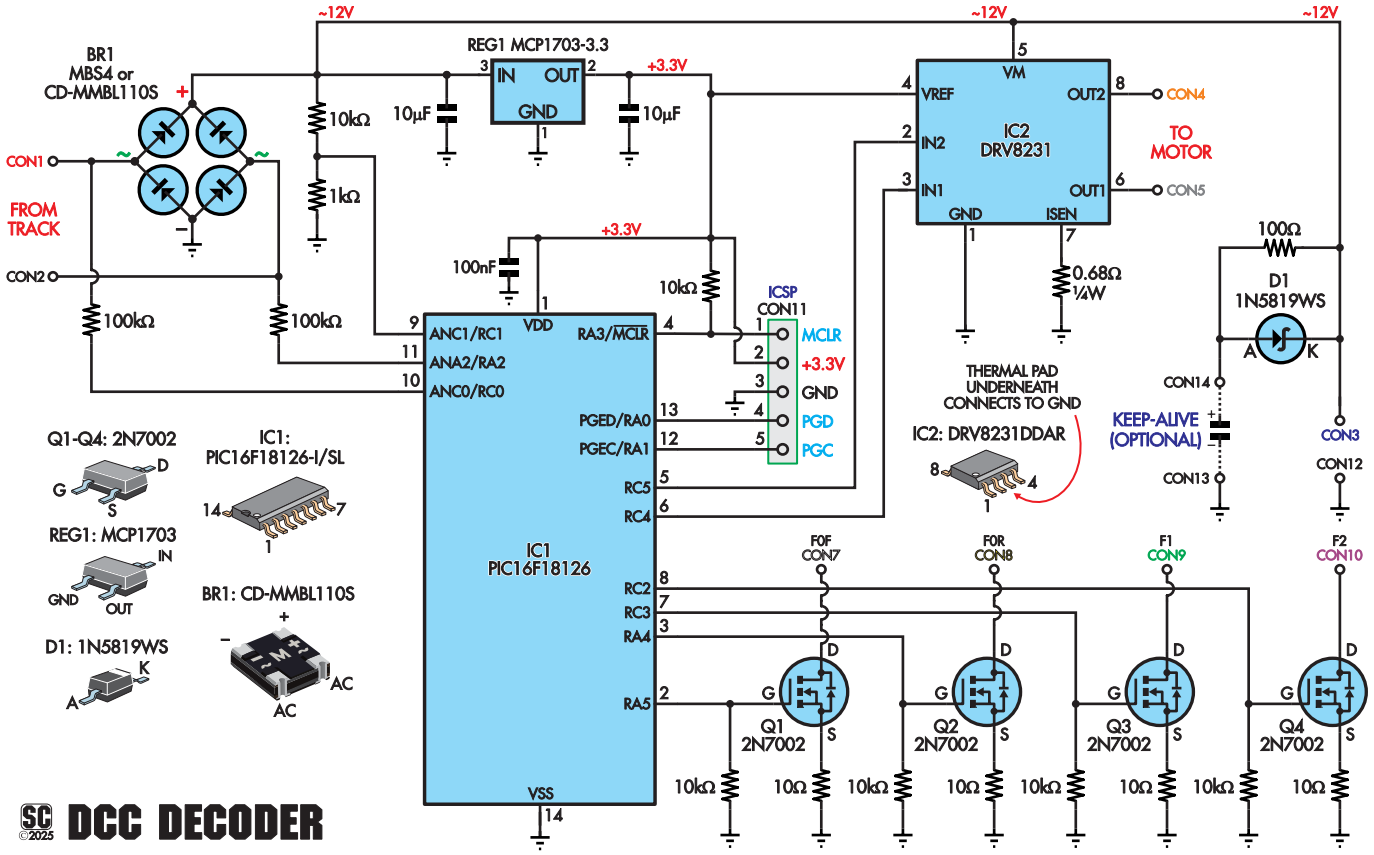

Fig.2: despite its simple design, the Decoder includes current limiting on its motor and function outputs. The 100Ω resistor provides charging current for an external capacitor that can feed power back to the circuit via D1, allowing the Decoder to continue operating through brief contact losses, as might occur with dirty track.

signals from IC1 (pins 5 and 6) feed control pins 2 and 3 on IC2. Pins 6 and 8 are the outputs, which are connected directly to the motor pads on the PCB (orange and grey); it is expected that a brushed DC motor is connected here.

IC2 sets our lower voltage operating limit. At about 4.5V, it disables its motor outputs, so we can guarantee operation down to 7V as required by the DCC standards. In practice, we expect that everything will work down to about 6V on the track.

The 3.3V rail also connects to pin 4 of IC2 (V_{REF}). This is used by the circuitry inside IC2 to provide a current-limiting feature. The motor current flows to ground via pin 7, so a 0.68Ω current measuring resistor (shunt) between pin 7 and ground develops a voltage proportional to the motor current.

An amplifier internal to IC1 multiplies this by 10 and compares it to the V_{REF} voltage. If V_{REF} has been exceeded, the power to the motors is shut off briefly to limit the current.

A current of 485mA develops 0.33V, sufficient to trip the current-limiting circuitry, hence the nominal 500mA motor current limit specification, which is well within the limits of this chip. The DRV8231 can also detect short circuits and, to limit damage, it provides a longer shut-off time in that case.

The circuitry around N-channel Mosfets Q1-Q4 is for the function outputs. Q1-Q4 are in an open-drain configuration for switching 12V loads under the control of a 3.3V microcontroller. The 10kΩ gate pull-down resistors ensure that the Mosfets are switched off any time the micro is not driving the control signals from its pins 2, 3, 7 and 8.

A typical load would be a light of some sort, such as a string of LEDs and their dropping resistors, or a small 'grain of wheat' incandescent globe, connected between the 12V pad (blue) and the function output (white, yellow, green or purple).

The 10Ω resistor between each Mosfet source and ground provides the current-limiting feature. The 2N7002 Mosfets require a source-to-gate voltage of about 2.3V to fully switch on, so if more than 100mA flows through one of the 10Ω resistors, there is no longer sufficient voltage (out of the 3.3V logic level signals) to turn on the Mosfet, and the current is limited to this level.

The Mosfet will not handle a short circuit to 12V indefinitely, even with current limiting, but will survive longer than if it were exposed to a direct short circuit. With four 100mA outputs, a 500mA motor plus some tens of milliamperes consumed by the microcontroller, the Decoder is also kept under the 1A limit of BR1.

The remaining circuitry is a so-called keep-alive that allows another larger capacitor to be connected, so that the Decoder can operate over dirty track or in other cases of intermittent DCC power. The capacitor is charged via the 100Ω resistor, limiting the peak charging current. The power is returned to the circuit when needed via D1, a schottky diode.

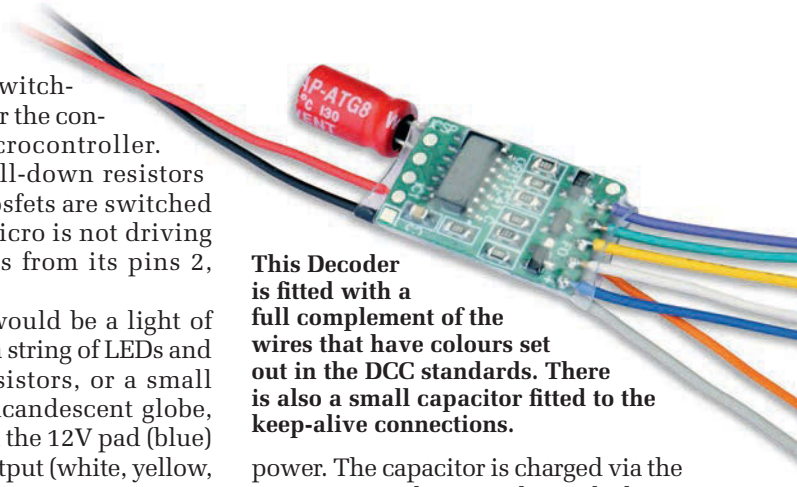

This Decoder is fitted with a full complement of the wires that have colours set out in the DCC standards. There is also a small capacitor fitted to the keep-alive connections.

We haven't specified a capacitor here, since the circuit is fully functional without it. If needed, the capacitor can be chosen based on the space available in the locomotive. It is completely optional and unnecessary in most cases. This is also a handy location to pick off 12V if needed for other circuitry.

The pins needed for in-circuit serial programming (1, 4, 12, 13 and 14) are connected to the ICSP header pads, with the standard 10kΩ pullup resistors between pins 4 (MCLR) and 1 (V_{DD}); pin 14 is ground. The 3.3V power and ground connections on the ICSP header can also be used to power other circuitry if required. Connection diagrams Figs. 3 & 4 show these locations on the PCB.

Firmware operation

The main role of decoding the incoming DCC signal is performed by an interrupt service routine (ISR) that is triggered every 22µs. DCC uses pulse widths of 58µs (nominal) to encode a binary '1' bit, with widths over 100µs encoding a '0' bit. By counting the time between DCC polarity changes, the bits can be decoded.

The Decoder then assembles the bits into packets, looking for a long string of 1s ('preamble') as the start marker. Since each byte is preceded by a zero bit, the firmware can easily determine its position within a packet.

Each packet also includes a checksum; if this doesn't match what the Decoder calculates, the packet is rejected as the data has been corrupted. Finally, the packet address (encoded as the first one or two bytes) is used to determine how the packet should be treated.

Parts List – DCC Decoder

1 double-sided 14 × 23mm PCB coded 09111241, 0.8mm thick
1 PIC16F18126-I/SL 8-bit microcontroller programmed with 0911124A.HEX, SOIC-14 (IC1)
1 DRV8231DDAR motor driver IC, SOIC-8 (IC2)
1 MCP1703A-3302 3.3V SOT-23 LDO regulator (REG1)
4 2N7002 SOT-23 N-channel Mosfets (Q1-Q4)
1 1A SMD bridge rectifier BR1 **[MBS4 or CD-MMBL110S]**
1 1N5819WS SOD-323 schottky diode (D1)
1 3cm length of 20mm diameter heatshrink tubing

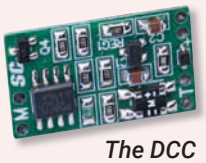

The DCC Decoder shown at actual size.

Capacitors *(all SMD M2012/0805 size MLCC)*
2 10µF 25V X5R
1 100nF 50V X7R

Resistors *(all SMD 1%, M2012/0805 size, ⅛W unless noted)*
2 100kΩ
6 10kΩ
1 1kΩ
1 100Ω
4 10Ω
1 0.68Ω ¼W

Kits for all designs in the DCC series, including this Decoder, are available to order from Silicon Chip's website: https://siliconchip.au/Shop/DCCKits

DCC Glossary

Accessory Decoder: A DCC decoder that is designed to be used in a fixed location; for example, to control points or signals, or fixed lights like streetlights. They operate within a separate address space to the decoders used in locomotives.

Analog: Describes the traditional mode of model railway operation, where a voltage is applied to the track that directly drives the locomotive's motor via wheel pickups. The analog controller sets the voltage and polarity to determine the speed and direction.

Base Station: A device that drives tracks with a DCC signal and has the means to accept user input to control the decoders. Many also incorporate programming features as well.

Booster: A device that can drive tracks with a DCC signal but has no user input. It might take its signal from a separate Base Station or other system. A booster is often used to increase the amount of power that can be delivered to the track.

Consist: In full-scale railway parlance, a consist is a grouping of rolling stock that operates together and might include locomotives, passenger cars and freight cars. In DCC, it describes a grouping of mobile decoders with a common consist address, allowing them to be operated as though they were a single unit.

CV (Configuration Variable): A non-volatile variable (typically stored in EEPROM) that determines a decoder's operating behaviour. This includes the address (fleet number) that it responds to and how its motors and functions react to user input.

DC: See Analog.

Decoder: A device that receives power and data from a DCC signal and drives lights, motors etc. Mobile decoders can be fitted to rolling stock, while accessory decoders can control trackside equipment.

Keep-Alive: A means of providing power to a mobile decoder if the track contact is poor or intermittent. It usually consists of a capacitor and some circuitry to keep the capacitor charged, feeding power back to the decoder as needed.

Locomotive: For DCC, this typically includes any item of rolling stock that is fitted with a motor and mobile decoder, including passenger railcars and motorised tenders.

Long Address: A 14-bit address with valid values from 0 to 10239. This range includes all four-digit numbers.

Mobile Decoder or Multifunction Decoder: Also known as a locomotive decoder, it is a decoder designed to operate the motor and lights of a mobile item of rolling stock, such as a locomotive, although it could be fitted to a passenger railcar or other powered unit.

Operations Mode: A form of CV programming that can occur on an operating track. The programming packets are addressed and can target a specific decoder. This is the only form of programming that can occur while the locomotive is active.

Programmer: A device that can perform CV programming. This might include a Base Station that includes programming features or a standalone device like our Arduino-based Programmer from October 2019.

Service Mode: Address-only, Physical, Paged and Direct modes do not program decoders by address, so they must occur on a segregated track in service mode to ensure that only the target decoder receives the programming commands.

Short Address: A 7-bit address in CV1. It is the only CV that is mandatory. Addresses up to 127 are valid, but values above 111 conflict with the addresses used for service mode programming packets and are best avoided.

As well as commands addressed to specific locomotives, there could also be broadcast packets on address 0. These allow all trains to be stopped quickly. Some packets are used to program configuration variables.

The 22µs timer is also divided down into 7ms intervals to manage the acceleration and deceleration of the motor. This is a feature of DCC that allows the locomotive speed to be ramped up or down at a rate to simulate inertia.

A separate 4Hz timer is used to allow the function outputs to be modulated in various flashing patterns. An eight-bit CV is used to set whether or not a function output should be switched on, so flash rates from ½Hz to 2Hz are possible, as well as other sequences.

Using these timers, the Decoder also monitors how long it has been since it has received a valid packet addressed to it. This can be used to stop the loco if a packet has not been received within a certain period.

If the firmware sees a correctly addressed packet, it is decoded and the internal state is updated. This could involve changing the speed or direction, or switching some of the function outputs on or off.

IC1's internal analog-to-digital converter is used to measure the voltage on the 12V rail. This can be used to compensate the motor and function outputs, to even out variations on the 12V rail. This can happen with dirty track or if the Decoder needs to source current from the keep-alive capacitor.

The motor driver is provided with a PWM (pulse-width modulated) signal to vary the motor speed. This raw speed value received is altered by the acceleration, deceleration and compensation features and can also be adjusted by means of CVs that can be used to alter the speed mapping profile.

There are also CVs that can be used to remap the function commands (as received in the DCC packets) to specific function outputs, so this is another processing step that occurs, as well as applying any flashing patterns that have been set.

CV programming

CV programming involves modifying the configuration variables (CVs) to change the Decoder's behaviour. The CVs are kept in non-volatile EEPROM, so are not lost during a power cycle. The list of CVs supported by our Decoder is shown in Table 1. These operate as described in the DCC standards, except for those set aside for manufacturer-specific features.

For example, CV1 is used to hold a 7-bit value that is known as the short address. Thus, changing CV1 allows the Decoder to respond to a different address. This is key to how DCC allows multiple locomotives to operate on the same tracks. Typically, the address will correspond to the fleet number painted on the locomotive or item of rolling stock.

The firmware also looks for the specific sequences needed for CV programming and updates the EEPROM and all live operating parameters. Since CV programming can make permanent changes, extra checks are made to validate these commands. One check that is performed (and required by the standard) is to only make changes when two consecutive, identical packets are received.

There are several different programming modes, but our Decoder fully supports Address-only, Physical, Paged and Direct modes. These fall into the

Table 1: configuration variables supported

CV	Notes	Default
1	A 7-bit value for the short address (the eighth bit is ignored). Values above 99 are not recommended because some of them conflict with the addresses used for service mode programming. Many systems restrict this to the range 0-99 and use the long address (CV17 and CV18) for all 3- and 4-digit addresses. The default address of 3 is standard.	3
2	Start voltage. This value (as a fraction of 255) dictates the proportion of voltage that is applied at the lowest speed setting.	0
3	Acceleration. A value of 0 means acceleration is disabled (the speed is updated instantly). Different systems are in use, but this Decoder applies a speed step every (7ms × value of CV3). Higher values mean slower acceleration, so a setting of 3 will result in one speed step every 21ms.	0
4	Deceleration. Works the same as Acceleration, but allows a different rate to be applied for deceleration.	0
5	High voltage. This value (as a fraction of 255) dictates the proportion of voltage that is applied at the highest speed setting. If it is 0 or 1, it is treated the same as 255, ie, full voltage.	0
6	Mid voltage. This value (as a fraction of 255) dictates the proportion of voltage that is applied at the middle speed setting. If it is 0 or 1, it is ignored (only the start and high voltages are used in calculations).	0
7	Manufacturer version number (read only).	0x5C ('SC')
8	Manufacturer identification number (read only). The value of 13 has been allocated by the NMRA for the identification of DIY decoders.	13
11	Packet timeout. This value dictates how many seconds a decoder will maintain speed without receiving a validly addressed packet. When set to 0, the decoder will not time out and will continue to maintain speed.	0
17	Most significant bits of the long address. The long address is nominally 14 bits and the top two bits of CV17 are set. The range of valid values for CV17 is restricted to 192-231, allowing the meaningful bits to represent the values 0-39. This gives an address range from 0-10239, which includes all four-digit decimal numbers.	192
18	The least significant bits of the long address.	0
19	Consist address and direction. Consisting allows multiple locomotives to be allocated to a train under the control of a single address, the consist address. The lower seven bits form a short address that can be used for consist addressing. When these seven bits are set to 0, consist mode is off; otherwise, the Decoder will respond to its consist address instead of its short or long address. When the upper bit is set, the locomotive operates in the opposite of its normal direction (eg, for running back-to-back with another locomotive).	0
29	Basic configuration. The bits in this CV each correspond to a basic configuration setting. Our Decoder implements: **Bit 0:** When set, the locomotive reverses its direction of travel. **Bit 1:** When set, the Decoder interprets baseline speed packets as 28-step; when clear, as 14-step. The Decoder will also respond to 128-step packets at all times. **Bit 5:** When set, the Decoder responds to its long address; otherwise, it uses its short address.	2
The default values for CVs 33-37 are 1, 2, 4, 8 and 0 respectively, meaning the functions are mapped across each row. It is possible to configure multiple commands to control a single output or for a command to control multiple outputs.		
33	Headlight forward command → F0F output (RA5)	1
34	Headlight reverse command → F0R output (RA4)	2
35	F1 command → F1 output (RC3)	4
36	F2 command → F2 output (RC2)	8
37	F3 command → Mapped to none	0
CVs in the range 47-64 are free for custom use by designers. We have implemented them as follows.		
47	Voltage compensation. When enabled, the Decoder adjusts the PWM duty cycle to counteract changes in the supply voltage. **Bits 0-3 (values 1-15):** when set, the Decoder sets the function outputs to match the voltage to the setting. For example, if the Decoder's supply voltage (measured at the bridge output) is 14V and bits 0-3 are set to 7 (0111), then the functions will be driven with a 50% duty cycle output. **Bit 4:** when set, the decoder compensates the motor PWM to make the full speed value correspond to 16V. When either value is set to 0, the corresponding compensation is disabled.	0
For CVs 49-52 each bit corresponds to a time slice of ¼ second over a two-second period, with the MSB occurring first. A set bit means the function output is on and thus different values will give different flash patterns. Where two different values are provided, they will operate out of phase. All values 0-255 are valid, see the right-hand column for what some example values will output.		***255: solid on 85 or 170: 2Hz 51 or 204: 1Hz 15 or 240: ½Hz***
49	F0F output (RA5) → Flash pattern	255
50	F0R output (RA4) → Flash pattern	255
51	F1 output (RC3) → Flash pattern	255
52	F2 output (RC2) → Flash pattern	255

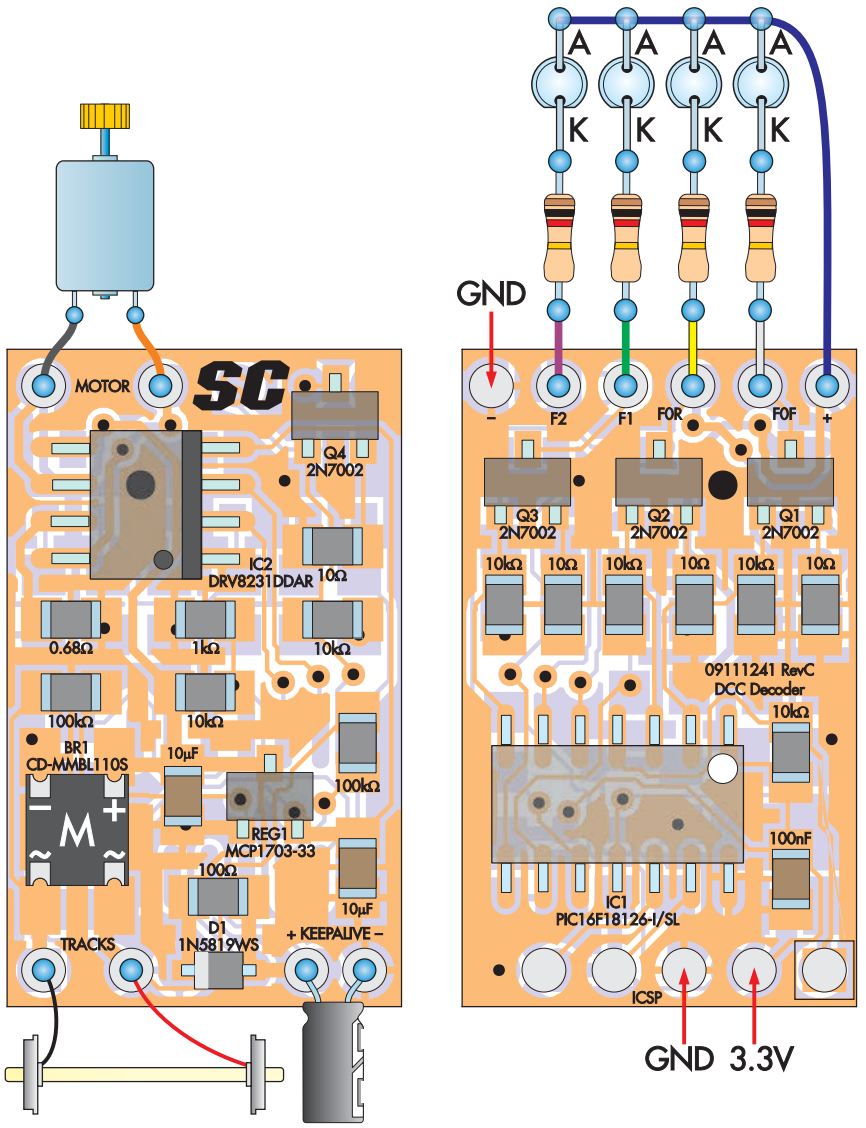

Figs.3 & 4: there is little room on the PCB for silkscreen markings, so you will have to pay close attention to these overlay diagrams, shown at ~400% scale. The Decoder will typically be wired into a locomotive as shown here. If using lamps, the polarity doesn't matter & series resistors are not needed.

category of what is known as service mode programming.

These modes do not use the Decoder address, so are typically done on a segregated programming track to target a single Decoder. They would most often be used during initial setup of the locomotive after the Decoder is installed, before it is placed on the mainline track.

Our Decoder also supports the long form of operations mode programming, which is intended to be used on a mainline track. It thus includes an address parameter, to ensure that the commands are directed to the correct decoder amongst the many that might be on the main track.

This mode could be used to tweak the speed or acceleration profile and then perform testing without having to shuffle the locomotive to and from the programming track. As such, this mode does not permit changes to critical CVs like addresses.

Decoder assembly

The Decoder is built on a double-sided PCB coded 09111241 that measures just 14 × 23mm. It has been deliberately kept as small as possible, to fit in tight spaces (it's even thinner than a normal PCB at 0.8mm), so there isn't much room for designators on the silkscreen. You will have to closely follow the overlay diagrams (Figs.3 & 4) to ensure all parts are fitted correctly.

The passive components are all M2012/0805 size (2.0 × 1.2mm), while the rest are in SOT-23 or SOIC packages. None are too difficult to solder, but it would be ideal for prospective constructors to have some SMD experience.

We recommend that you have on hand flux paste (ideally in a syringe), solder-wicking braid, a magnifier, fine-tipped tweezers and the usual SMD gear. We found it easier to start assembly with the top side of the PCB; this has the track and motor connections (M and T) and the SOIC-8 chip, IC2. A PCB holder or some Blu-Tack will be handy to secure the board and prevent it from moving around.

IC2 has an exposed pad on its underside. We have placed a large hole under that pad so that you can flow solder through from the other side, or melt solder paste via that hole.

To ensure that solder will flow through this hole from the other side of the board, add plenty of flux paste to the pad under the IC and inside that hole. That should help to draw solder through from the other side of the board.

Solder this first so that you can check that this joint is secure before starting on any other pins. This could also be soldered with a reflow oven, if you have one. We secured the chip flat against the PCB with polyimide (Kapton) tape and then flipped the board over to feed solder through the pad hole on the other side.

Flow solder into the hole and apply heat for a few seconds. If the chip is properly soldered, after removing the tape, the chip remains in place. Then solder the remaining leads. After that, fit REG1 as shown. The single diode, D1, has its cathode towards the bridge rectifier and track connections.

The two capacitors on this side are both 10µF parts. Once they are in place, solder the bridge rectifier, being sure to orientate it correctly. Follow with Q4 and the resistors on this side, being careful to match their values. Flip the PCB over and solder Mosfets Q1-Q3 and IC1. Note that both ICs have their pin 1s towards the centre of the PCB.

The sole 100nF capacitor is in one corner of the PCB. The seven remaining resistors can then be fitted. Clean off the excess flux with an appropriate solvent and allow the PCB to dry, inspect it carefully and rectify any solder

We have wired up this compact N-scale motorised chassis (Kato 11-107) using our Decoder, allowing us to run tests on a short length of track. We removed two folded metal strips from inside the chassis to break the electrical connection between the track pickups and the motor. For clarity, we have left off the heatshrink, but we recommend you fit it to prevent short circuits. For views inside the mechanism, see www.hookstonemodels.co.uk/dcc-conversion/kato/11-107-chassis

bridges or dry joints before powering up the Decoder.

Programming IC1

If you have purchased a kit or programmed IC, IC1 will be programmed and programming will not be needed. Otherwise, you can use a Snap, PICkit BASIC, PICkit 4 or PICkit 5 to program the 0911124A.HEX file onto the chip.

We were able to program one of our prototypes by fitting a 5-way pin header strip to a PICkit 5's header socket and then pressing that against the Decoder's ICSP header pads. We used the PICkit 5 to provide 3V power.

If that doesn't work, then you may need to temporarily solder a header strip to the Decoder's ICSP header pads. The pad spacings are slightly narrower than 0.1in (2.54mm), but a standard 0.1in header can be made to fit without too much trouble.

Testing

If you would like to run some tests on the Decoder, you can apply power from a 9V battery to the 'T' pads that feed into the bridge rectifier. Because of the bridge, the polarity does not matter.

This should power up the Decoder and you should see about 7.5V across the keep-alive pads. The output of the 3.3V regulator should be available on pins 2 and 3 of the ICSP header. You can see these connections in Figs. 3 & 4. If these voltages are present and nothing is getting hot, all is as well as can be expected.

Connections

The connections to the Decoder are shown in Figs.3 & 4, with the wire colours chosen to match those set out in the DCC standards. The other connections are not part of the standard, but could be useful. If you have previously installed a decoder in a locomotive, these colours should be familiar.

There is a lot of detail in fitting a Decoder to a locomotive. We can only provide some general advice, but anyone with a reasonable understanding of electronics should be able to figure it out.

It's a good idea to check the voltage and current demands of the locomotive and any additions before making any changes. We expect that LEDs and most small, modern motors will satisfy the requirements, but incandescent globes and older motors may draw more current than the Decoder can provide.

One of the most important things is to ensure that all existing connections between the wheel pickups (to the track) and the motor or lights are broken. An existing connection may conflict with a connection made by the Decoder, causing a short circuit.

In Fig.1(a), the red lines show the wires that might typically need to be cut; note that the diodes will probably be discarded in this case, as the Decoder provides rectified DC to the function outputs. Check the polarity of any LEDs carefully as well, since they may need to be reversed.

In practice, the red and black (track) wires will go to the wheel pickups on the locomotive. They are effectively interchangeable, but the standards specify that the red wire should connect to the right-hand rail when the locomotive is travelling in the forward direction.

The grey and orange wires go to the motor and, again, they are interchangeable in practice. The standards say that it should be wired up so that the locomotive moves forward when the orange wire has the more positive voltage.

We suggest you connect these wires in whatever way is easiest. If it turns out that the locomotive drives the wrong way, that can be fixed by simply changing a configuration variable.

The blue wire effectively provides a positive supply, with the various function outputs connecting to negative when active. Thus, LEDs should have their anodes connected to the blue wire; common-anode LEDs are typical. There is usually enough voltage to allow two or maybe three LEDs to be connected in series with a single ballast resistor.

The default configuration makes the white wire active when the locomotive is going forward and the yellow

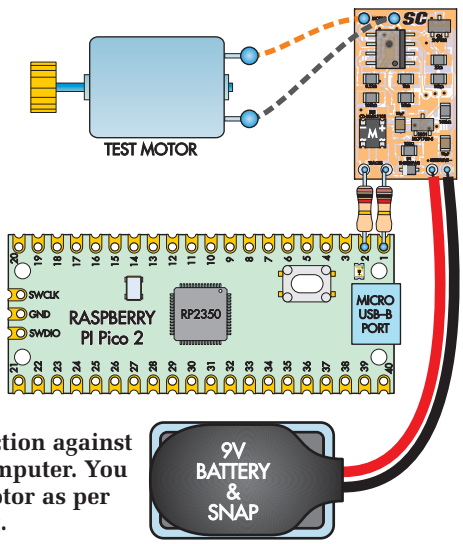

Fig.5: wire up a Pico 2 and a 9V battery like this to use the Arduino sketch to test the Decoder. The resistors are simply provided for protection against damage to the Pico 2 or a connected computer. You should also connect some lights or a motor as per this diagram to see the outputs working.

DOWNLOAD PE BACK ISSUES

COMPLETE YEAR collections
From just £5.95 each

VOL 28 – January to December 1999 – £5.95
VOL 29 – January to December 2000 – £6.95
VOL 30 – January to December 2001 – £7.95
VOL 31 – January to December 2002 – £7.95
VOL 32 – January to December 2003 – £7.95
VOL 33 – January to December 2004 – £7.95
VOL 34 – January to December 2005 – £8.95
VOL 35 – January to December 2006 – £8.95
VOL 36 – January to December 2007 – £8.95
VOL 37 – January to December 2008 – £8.95
VOL 38 – January to December 2009 – £10.95
VOL 39 – January to December 2010 – £11.95
VOL 40 – January to December 2011 – £12.95
VOL 41 – January to December 2012 – £14.95
VOL 42 – January to December 2013 – £15.95
VOL 43 – January to December 2014 – £16.95
VOL 44 – January to December 2015 – £17.95
VOL 45 – January to December 2016 – £19.95
VOL 46 – January to December 2017 – £22.95
VOL 47 – January to December 2018 – £27.95
VOL 48 – January to December 2019 – £27.95
VOL 49 – January to December 2020 – £29.95
VOL 50 – January to December 2021 – £32.95
VOL 51 – January to December 2022 – £34.95
VOL 52 – January to December 2023 – £34.95
VOL 53 – January to December 2024 – £34.95

FIVE-YEAR collections – 60 issues!
From just £14.95 each

PE2832: Jan 99 to Dec 03 – £14.95
PE2933: Jan 00 to Dec 04 – £14.95
PE3034: Jan 01 to Dec 05 – £14.95
PE3135: Jan 02 to Dec 06 – £14.95
PE3236: Jan 03 to Dec 07 – £14.95
PE3337: Jan 04 to Dec 08 – £14.95
PE3438: Jan 05 to Dec 09 – £19.95
PE3539: Jan 06 to Dec 10 – £25.95
PE3640: Jan 07 to Dec 11 – £25.95
PE3741: Jan 08 to Dec 12 – £25.95
PE3842: Jan 09 to Dec 13 – £33.95
PE3943: Jan 10 to Dec 14 – £33.95
PE4044: Jan 11 to Dec 15 – £33.95
PE4145: Jan 12 to Dec 16 – £35.95
PE4246: Jan 13 to Dec 17 – £37.95
PE4347: Jan 14 to Dec 18 – £38.95
PE4448: Jan 15 to Dec 19 – £39.95
PE4549: Jan 16 to Dec 20 – £42.95
PE4650: Jan 17 to Dec 21 – £44.95
PE4751: Jan 18 to Dec 22 – £44.95
PE4852: Jan 19 to Dec 23 – £44.95
PE4953: Jan 20 to Dec 24 – £44.95

20 YEARS OF PE!
AMAZING BUNDLE OFFER
Every issue from 1999 – 2018
240 magazines for just

£69.95

PE2832: Jan 99 to Dec 03 PLUS!
PE3337: Jan 04 to Dec 08 PLUS!
PE3842: Jan 09 to Dec 13 PLUS!
PE4347: Jan 14 to Dec 18

Download from:
www.electronpublishing.com

wire active when the locomotive is in reverse. Again, it might make sense to use whatever wire is convenient and change the function mapping using the configuration variables.

Remember to thread heatshrink tubing over the wires before soldering them, then shrink it into place once everything is tested and operational to prevent accidental short circuits.

If you don't plan on wiring an external capacitor to the keep-alive pads, you might be able to fit another surface-mounting capacitor (similar to the 10µF parts already on the PCB) directly across these pads. That will give the Decoder a small extra power reserve.

Arduino test sketch

We are providing an Arduino test sketch (DCC_ADD3_TEST) that can be used if you have a locomotive (or motors and lamps) and Decoder, but not a base station. It is written for a Pico or Pico 2 using the arduino-pico board profile. It is based on the code we have developed for the upcoming Base Station.

The sketch produces a logic-level DCC signal between GP0 & GP1 pins and expects power to be provided by other means. We can use the keep-alive connections for this purpose; something like a 9V battery would work and provide a limited current. Note that the polarity is important when connecting to the keep-alive pads.

Fig.5 shows the wiring needed to run a test with the Arduino sketch. You will also need to connect a motor or lights to see any effect. The sketch produces commands with the default address of 3, and drives the motor outputs in both the forwards and reverse directions. The function outputs are activated in turn, with the actions shown on the serial monitor.

Conclusion

Readers who already have a DCC system will quickly integrate DCC Decoders into their layout and have their own ideas about how to set them up.

A DCC Decoder is not much fun on its own, so we will be following up with the DCC Base Station next month. We will also provide more background with a feature article for getting started with DCC in case you have not worked with a DCC system before. PE

Subscribe to SILICON CHIP

SILICON CHIP is one of the best DIY electronics magazines in the world. Each month is filled with a variety of projects that you can build yourself, along with features on a wide range of topics from in-depth electronics articles to general tech overviews.

	Print	Combined	Online
6 months	$115 ≈£60 ≈€70	$125 ≈£65 ≈€76	$55 ≈£29 ≈€34
1 year	$220 ≈£115 ≈€134	$240 ≈£125 ≈€146	$105 ≈£55 ≈€64
2 years	$410 ≈£214 ≈€250	$445 ≈£233 ≈€271	$200 ≈£105 ≈€122

- *Combined subscriptions include both the printed magazine and online access.*
- *Prices are in Australian dollars (AUD) unless otherwise specified.*
- *Prices listed are for everywhere in the world except Australia and New Zealand. View siliconchip.com.au/Shop/SubRates for a full list of current prices.*
- *Price estimates are just that and should only be used as a guide. We can only take payment in the form of Australian dollars.*

Published in SILICON CHIP

Destination Display; August 2026

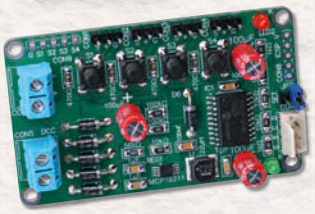

DCC Accessory Decoders; July 2026

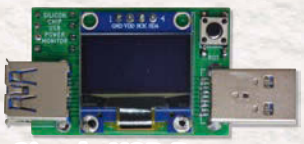

Simple USB Power Monitor; June 2026

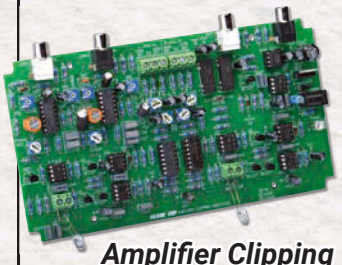

Amplifier Clipping Indicator; May 2026

Try our Online Subscription – now with PDF downloads!

An online issue is perfect for those who don't want too much clutter around the house and is the same price worldwide. Issues can be viewed online, or downloaded as a PDF.

To start your subscription go to

siliconchip.com.au/Shop/Subscribe

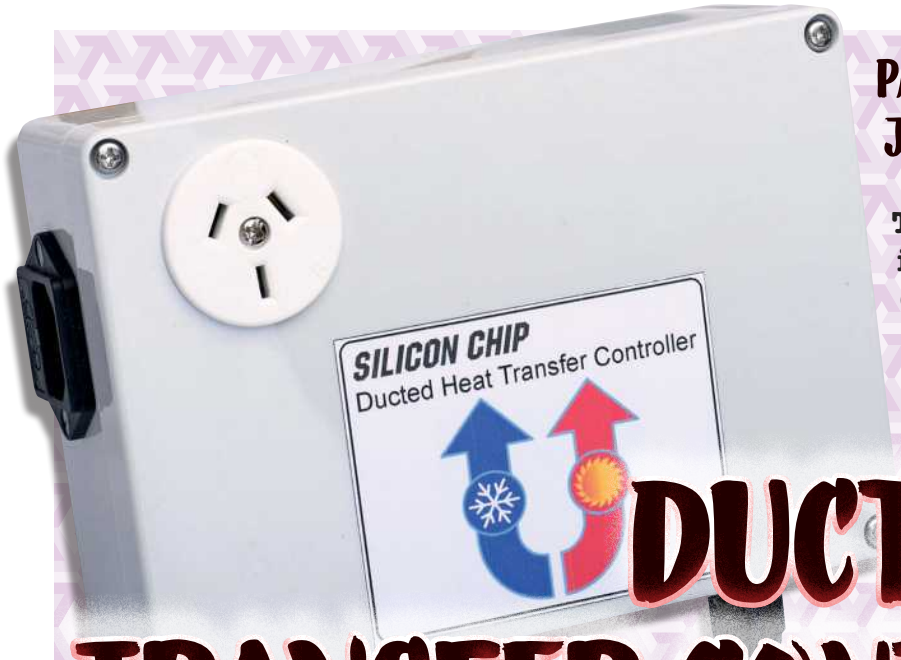

PART 1 BY
JULIAN EDGAR & JOHN CLARKE

This smart controller can improve the energy efficiency of your home. It can transfer warm or cool air between rooms automatically when needed.

DUCTED HEAT TRANSFER CONTROLLER

This device controls a mains-powered fan that is used to transfer heat between rooms via ducts. The controller can be used manually, automatically, or based on a timer. The wall-mounted LED gives an indication of the temperature difference between rooms.

Advantages

The most common reason for using a ducted heat transfer system is when the source of heat for the house is confined largely to one room. There are two likely situations where that would occur: a wood heating stove is located in one room, or passive solar heating occurs largely at one end of the house.

While in some jurisdictions, wood heating is frowned upon (for example, the Australian Capital Territory is phasing out wood heaters), wood heaters can be environmentally acceptable and, in some areas, cost little to run. Wood heaters are effectively carbon neutral; the carbon dioxide absorbed by the trees during their growth is released when the wood is burnt.

Wood heaters have a bad reputation for emissions – we've all seen wood heater flues emitting a stream of smoke for many hours. There are two reasons reason for this. The first is that the heater has been shut down – the damper closed to reduce the airflow. The second reason for smoke emis- sions is burning green wood that has high moisture levels .

With current heaters that must meet emissions standards, a wood heater burning dry wood at full power produces no visible emissions. But the key point is 'at full power' – throttling the heater output reduces its efficiency and increases emissions.

That's where a fan-forced heat trans- fer duct comes in. It is much better to keep the wood heater burning furi- ously and transfer some of that heat to other rooms in the house than it is to shut the heater down. All New Magazines softarchive.is Since most homes using wood heating have only one heater, using a transfer duct also works to warm more than just the room where the heater is located (see Fig.1).

The second reason for using a ducted heat transfer system is in houses that use passive solar heating. In southern Australia, windows facing north can be used to warm the house in winter. The sun shines in through these win- dows, heating the wall and floor sur- faces of the room, and subsequently the air within.

Because the sun is higher in the sky in summer, projecting eaves can shade these windows in summer, so

Features & Specifications

- Powered by the 230V AC mains
- Operates during all seasons without changes
- Three different operating modes
- Adjustable temperature difference and hysteresis
- Optional adjustable timer
- Optional fire alarm feature
- Wall plate button with sound and LED indicators
- Sensor disconnection indication
- **Temperature difference options:** 1, 1.5, 2, 3, 4, 5, 6, 8, 10 or 11°C
- **Hysteresis options:** 0.5, 1, 1.5, 2, 3, 4, 5, 6, 8 or 10°C
- **Timer options:** 15m, 30m, 1h, 2h, 3h, 4h, 5h, 6h, 8h, 12h or multiples thereof
- **Modes:** manual, timed, or automatic
- **Fire alarm function:** switches on RLY2 and rapidly pulses the piezo buzzer and LED when the temperature rise of either sensor is >8°C per minute or 70°C is exceeded (this does not replace a smoke alarm!)
- **Maximum total fan current:** 10A

Ducted Heat Transfer Controller

Air return paths are required

A heat transfer duct works by moving air – that is, pushing air from one room to another. But unless the air has a return path, the duct will not be very effective. Without a return path, air pressure will rise in the destination room, slowing the transfer of air. It's therefore best to leave some internal doors open so that good circulation can be achieved.

WARNING: Mains Voltage

This Direct Heat Transfer Controller operates directly from the 230V AC mains supply; contact with any live component is potentially lethal. Do not build it unless you are experienced working with mains voltages.

the northern windows don't heat the house when you don't want them to. In the northern hemisphere, this is reversed – you want southerly windows.

However, the number of rooms in a house that can face north is quite limited, so this type of heating can usually work in only one or two rooms. That's especially the case if the house was never designed with passive solar heating in mind. In this case, a ducted heat transfer system can be added to move solar heat to other rooms.

The problem with commercial options

Fan-forced heat transfer ducts are commercially available for installation in new or existing builds (see the photo overleaf). Typically, they comprise flexible ducting and one or two mains-powered fans. Common duct and fan diameters are 150mm, 200mm, 250mm and 300mm. The fan and duct are usually mounted in the roof space with the inlet and outlet grilles located in the ceiling.

Generally, these require you to switch them on manually when desired. You can certainly do that, but it's a little trickier than it first appears.

One thing that makes it tricky is that the temperature differences can be very small. For example, in a house that uses passive solar heating, the temperature difference from the 'warm' northern room to the southern 'cool' part of the house may initially be only 2°C. That difference may increase quite slowly – over hours.

Without either walking back and forth to feel the temperatures, or consulting room thermometers, the best time to turn on the fan isn't at all obvious. That's if you're even home at the time!

Luckily, this Transfer Controller can do the work for you.

Also, you may want the fan to operate for some time after you go to bed – you're no longer in the heated lounge room, and you want that residual heat distributed through the house. Or you want to be manually in charge of when the fan operates, but with a monitoring LED showing when the heated room is warmer by, say, 3°C than the room at the other end of the duct.

Our Controller can perform all these functions.

In long ducts, more than one fan may be needed. The controller can run fans up to a total power consumption of 2300W (10A at 230V). Since most duct fans are quite low in power, it can likely drive however many fans you need. If running multiple fans in the duct, ensure they both blow in the same direction!

Operating modes

The main function of the Ducted Heat Transfer Controller is to switch on the fan in the duct – the output is simply on or off.

However, when it activates that fan depends on the mode. Each mode is selected by switch BCD4 on the printed circuit board (PCB) – as with the other set-up features, it is expected that this will be set and then not frequently changed.

In all modes, the user interface is a neat wall-mounted, spring-return rocker-type pushbutton with a white monitoring LED visible around the periphery of the button, and a beeper mounted behind. The other two inputs are temperature sensors – one in the room at each end of the duct.

Mode 0 is manual mode. In this case, the pushbutton is used to switch the fan on and off.

Mode 1 provides manually triggered timed operation. Pressing the pushbutton switches the fan on for a specified period. Each quick press of the button adds (for example) one hour of operation, so one press gives one hour, two presses

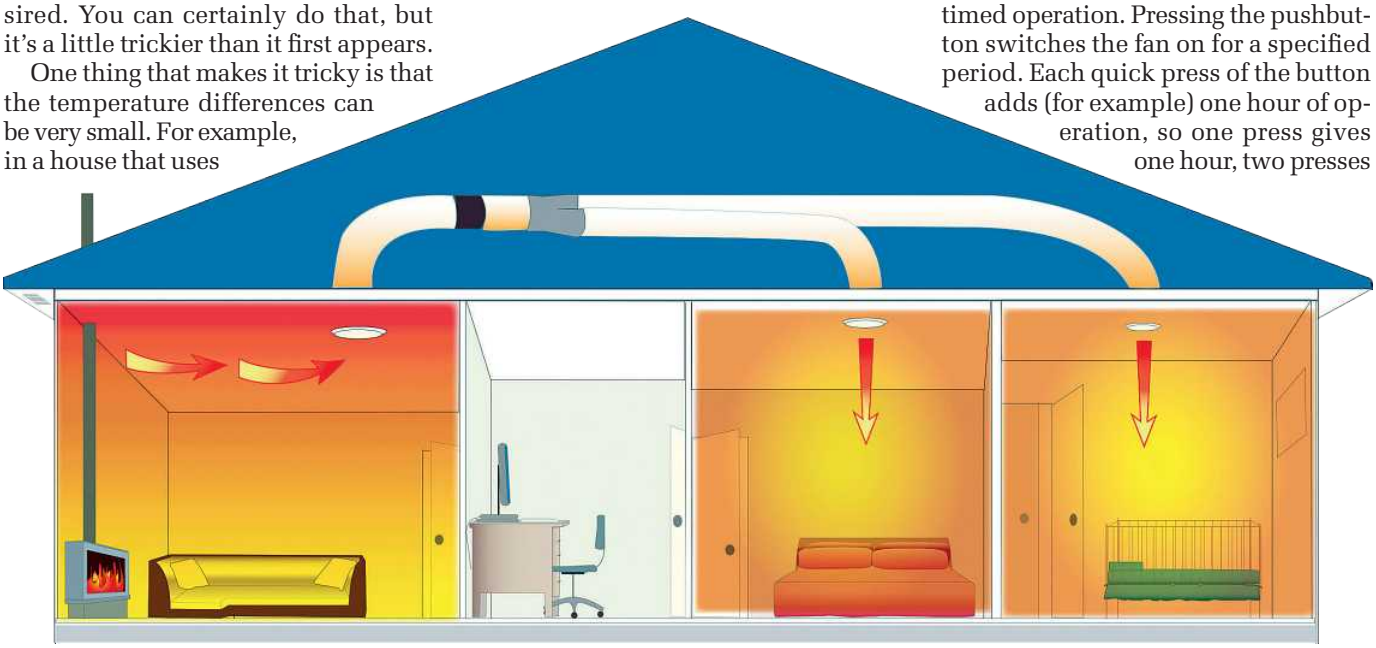

Fig.1: a Ducted Heat Transfer System takes the heat from one room and distributes it to one or more other rooms. A fan in the duct is used to move the air, and our controller determines when the fan switches on. Source: Vent-Axia.

This 150mm Ducted Heat Transfer System uses a single fan to distribute the air to two other rooms. Note that this duct is uninsulated – not a good idea.
Source: JPM Brands

Table 1 – BCD switch settings

Switch	BCD1 (temp. difference)	BCD2 (hysteresis)	BCD3 (timer period)	BCD4 (mode)
0	1°C	0.5°C	15 minutes	Manual
1	1.5°C	1.0°C	30 minutes	Timed
2	2°C	1.5°C	1 hour	Automatic
3	3°C	2°C	2 hours	Automatic
4	4°C	3°C	3 hours	Automatic
5	5°C	4°C	4 hours	Automatic
6	6°C	5°C	5 hours	Automatic
7	8°C	6°C	6 hours	Automatic
8	10°C	8°C	8 hours	Automatic
9	11°C	10°C	12 hours	Automatic

What about transferring cool air?

While we have described the function of the controller as operating a fan-forced duct that transfers warm air to a cooler room, the system can also transfer cool air to a warmer room. In fact, no changes are needed to do this because the system operates based on the temperature difference between the two rooms, rather than how much cooler the room is at the far end of the duct.

For example, say you have the difference in room temperature set to 3°C and the Mode set to 2 (Automatic). When the room at the end of the duct is 3°C warmer than the room at the beginning of the duct, the fan will switch on, transferring cooler air to the hotter room. Of course, the source room needs to be the same room in both winter and summer.

gives two hours etc, up to a maximum of five presses. A BCD switch preset determines the base period, from 15 minutes to 12 hours.

Mode 2 is fully automatic. In this mode, the fan operates when the temperature difference between the two ends of the duct exceeds a preset threshold.

In addition to mode selector switch BCD4, the PCB has three more adjustments.

BCD1 is used to set the temperature difference that needs to occur before the fully automatic mode (Mode 2) switches on the fan. This can be set to 1, 1.5, 2, 3, 4, 5, 6, 8, 10 or 11°C.

BCD2 is used to set the hysteresis. This is the difference between the switch-on and switch-off temperatures. This can be set to 0.5, 1, 1.5, 2, 3, 4, 5, 6, 8 or 10°C. It must be set lower than the temperature difference.

Let's imagine the temperature difference is set to 4°C and the hysteresis is set to 1°C. If the heated room is at 20°C and the unheated room is at 16°C (a difference of exactly 4°C), the fan will switch on. It will stay on until the difference in temperature decreases to 3°C; eg, the unheated room warms to 17°C.

In use, if the fan switches on and off too frequently, increase the hysteresis setting. On the other hand, if the temperature of the room at the other end of the duct varies up and down too much, decrease the hysteresis.

BCD3 sets the timed period that occurs in Mode 1 with each button press. In the example above, I suggested that each press gives a one-hour extension of the on-time. However, each button press can actually be set to be 15m, 30m, 1h, 2h, 3h, 4h, 5h, 6h, 8h or 12h.

Refer to Table 1 for all the BCD switch settings.

Monitoring LED and beeper

In addition to the pushbutton switch, the wall-mounted indicator is equipped with one LED and a beeper. The beeper operates in the same way in all modes: a single beep indicates switch on (a short press of the button) and a triple beep indicates switch off (achieved by a longer press of the button). The triple beep comprises a single beep followed by a quick double beep.

The LED can show different information in each mode.

In manual mode, if the fan is off, the LED is off, possibly flashing on briefly.

If the fan is on, the LED is on, possibly flashing off briefly.

If it's flashing briefly every two seconds when the fan is off, that indicates the measured temperature difference is greater than or equal to the set temperature difference, so you might want to switch it on. Similarly, if it's briefly flickering off while the fan is on, that means the temperature difference has fallen below the set difference (including hysteresis), indicating you may want to switch it off.

Manual Timed operation (BCD4 position 1) has LED behaviour that is the same as the manual mode.

Automatic mode (BCD4 position 2) has different LED behaviour. If the system has been disabled, the LED flashes. If the fan is on, so is the LED. If the fan is off, again, so is the LED. A summary of these modes is shown in Table 2.

Other potential uses

This device can also control a powered ventilator or fan; for example, one that ventilates a hot roof cavity in summer. In this use, one temperature sensor is placed in the roof cavity (or other hot area needing ventilation) and the other outside in an area protected from the weather (eg, under the eaves).

In this application, the best settings will probably be Mode 2 (automatic), with the temperature difference set higher than you would use for internal house use (eg, 10°C with 5°C of hysteresis).

Another use is for solar air heaters. While uncommon in Australia, these have been widely used in solar homes in the United States. In this approach, air is heated by a flat plate collector – a little like a traditional solar water heater but with air rather than water heated through contact with the plate.

When the air in the heater rises sufficiently in temperature, a fan can be used to move that heated air into the house through conventional air-conditioning ducts. In this application, one sensor would be placed so that it is exposed to the air in the heater (but shielded from direct sunlight), while the other would be placed inside the house. The temperature difference would be set quite low (eg, 3-4°C, with perhaps 2°C of hysteresis).

Circuit details

The complete circuit for the Ducted Heat Transfer Controller is shown in

Parts List – Ducted Heat Transfer Controller

1 Gainta G317-IP67 or G218 222 × 146 × 55mm ABS or polycarbonate IP65/IP67 enclosure [TME G317, TME G218]
1 double-sided, plated-through PCB coded 17101251, 151 × 112mm
1 lid panel label (84 × 65mm) and side panel label (64 × 10.5mm)
1 3VA 9+9V PCB-mounting mains transformer (T1) **[eg, BLOCK VB3.2/2/9 from Farnell]**
1 FRA4 or equivalent 250V 30A AC SPST relay with 12V DC coil (RLY1) **[eg, HF105F-4/012D1HSTF]**
1 NT73-2 PCB-mounting 250V 10A AC SPDT relay with 12V DC coil (RLY2)
4 PCB-mounting 10-position BCD switches (BCD1-BCD4) ***OR***
4 2×4-pin headers and 12 jumper shunts
1 2-way header, 2.54mm pitch (JP1)
1 jumper shunt (JP1)
2 15A 300V 2-way screw terminals, 8.25mm pitch (CON1, CON2)
1 2-way screw terminal, 5/5.08mm pitch (CON3)
1 3-way screw terminal, 5/5.08mm pitch (CON4)
3 8P8C RJ45 PCB-mounting horizontal sockets (CON5-CON7) **[eg, Adam Tech MTJ-881X1]**
1 IEC C14 panel-mount mains input connector
1 mains lead with IEC plug
1 13A chassis-mount mains socket [eg, RS 500-0459]
1 20-pin DIL IC socket (**optional**, for IC1)
1 fast-blow 10A M205 fuse (F1)

Hardware
2 M4 × 10mm panhead machine screws with matching hex nuts
2 M3 × 15mm panhead nylon machine screws
5 M3 × 6mm panhead machine screws
3 M3 brass hex nuts
1 200mm cable tie and 8 100mm cable ties
1 3-6.5mm diameter wire entry cable gland

Value	4-band code	5-band code
470Ω		
1.0kΩ		
2.2kΩ		
10kΩ		

Wire & cable
1 200mm length of black 7.5A hookup wire
1 50mm length of light-duty red hookup wire and light-duty black hookup wire
assorted lengths of 10A mains-rated green/yellow striped wire (150mm length); brown wire (200mm length); and blue wire (100mm length)
3 Cat 5, Cat 5E or Cat 6 patch leads, lengths to suit installation
assorted lengths of clear heatshrink tubing (70mm length, 5mm diameter; 30mm length, 4mm diameter; and 50mm length, 1mm diameter)

Semiconductors
1 PIC16F1459-I/P microcontroller programmed with 1710125A.HEX, DIP-20 (IC1)
1 7805 1A 5V linear regulator, TO-220 (REG1)
3 BC337 NPN transistors, TO-92 (Q1-Q3)
1 W02(M) or W04(M) 1.5A 200V/400V bridge rectifier (BR1)
16 1N4148 200mA 75V diodes (D1-D16)
3 1N4004 1A 400V diodes (D17-D19)

Capacitors *(16V PC radial electrolytic, unless specified)*
2 470µF 1 100µF 2 100nF 63V/100V MKT polyester

Resistors *(all ¼W, 1%)*
5 10kΩ 2 2.2kΩ 4 1kΩ 1 470Ω

Control panel parts (per panel)
1 double-sided, plated-through PCB coded 17101253, 51 × 67mm
1 Click MiniGrid CMA401 Single Gang Unfurnished Wall Plate ●
1 Click MiniGrid MD004WH 2-Way Retractive Switch Module ●
1 panel label, 45 × 30.5mm
1 top-entry 8P8C vertical RJ45 socket (CON10) **[eg, MOLEX 95503-2881]**
1 3-16V self-oscillating PCB-mount piezo buzzer, 7.5-7.6mm pitch, ≤14mm diameter
1 2-way vertical polarised header, 2.54mm pitch, with matching plug and pins (CON11)
1 2-way terminal block, 5/5.08mm pitch (CON12)
1 8P8C double adaptor (only required if using two control panels)

Temperature sensor parts
2 60 × 60 × 20mm vented enclosures or similar **[eg, Hammond 1551V3BK or 1551V3GY]**
2 double-sided, plated-through PCBs coded 17101252, 20 × 37.5mm
2 8P8C RJ45 PCB-mounting horizontal sockets (CON8, CON9) **[eg, Adam Tech MTJ-881X1]**
2 DS18B20 temperature sensors (TS1, TS2)

● *available from electrical wholesalers, such as Scolmore and Click4Electrics.*

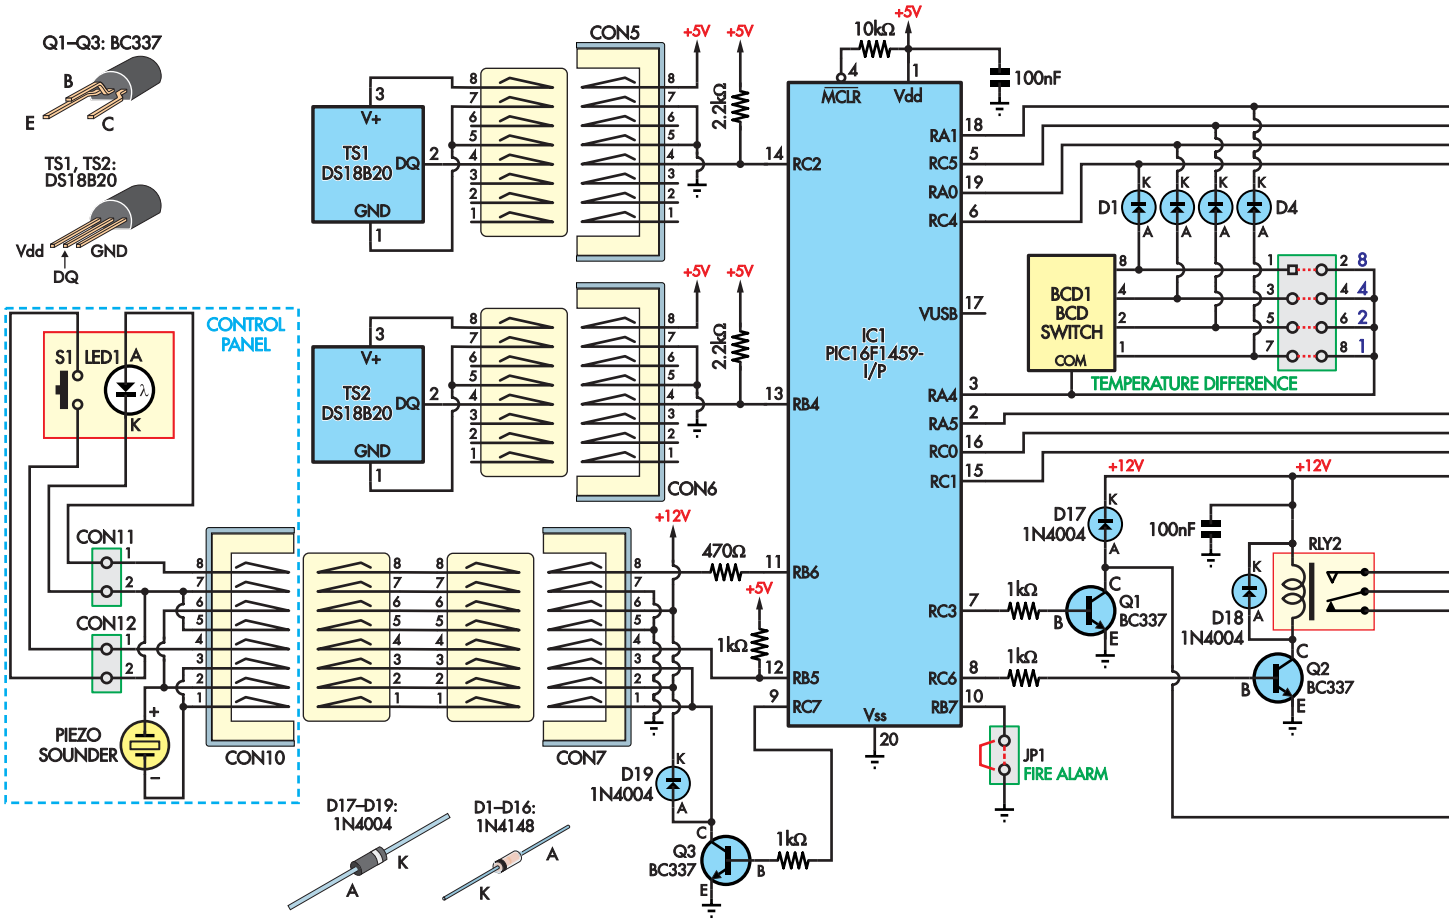

Fig.2. Microcontroller IC1 monitors the temperatures via sensors TS1 & TS2, which connect to the main board via 8-way Cat 5 (or similar) cables and RJ45 plugs/sockets. In each case, pin 4 carries the digital signal, pin 8 the 5V supply for the sensor and pins 5 & 7 are grounds.

TS1 & TS2 are Maxim DS18B20 1-wire digital thermometers. Just one data line (DQ) is required for serial communications. A minimum of one extra connection for the common ground connection is also required. Power for the sensor can be derived from the data line, but we include a direct 5V supply connection (V_{DD}/V+) since we have enough wires and this makes signalling easier.

Two-way communication between the microcontroller and temperature sensor is possible since the DQ pin is an open drain with a pull-up resistance of 2.2kΩ. Open drain means that the drain of a Mosfet connects to this pin, so when the Mosfet is on, the pin is pulled to 0V, while if it is off, it is pulled up by the resistor.

A Mosfet at either end of the wire can be used to pull it down to 0V, so a signal can be sent by the device at either end of the wire. The microcontroller uses its RC2 and RB4 I/O pins to request temperature readings and get them from the sensors.

The DS18B20 has a temperature reading accuracy of ±0.5°C from -10°C to +85°C. Temperature readings are available in 0.125°C steps, but for this project, we measure the temperature in 0.5°C increments.

Enabling the fire alarm feature

The Ducted Heat Transfer Controller can also be configured as a fire alarm. Because the system has temperature sensors that would normally be placed at divergent ends of the house, monitoring of these sensors provides a widespread back-up system to the legally required smoke detectors.

When this function is enabled by shorting the pins of JP1, each temperature sensor is monitored for both the temperature and the rate of temperature change. If the temperature exceeds 70°C and/or the rate of temperature change exceeds 8°C per minute, the beeper and LED rapidly pulse. Relay RLY2 is also energised, which can power a low-voltage warning siren, switch on low-voltage lights etc.

If the fire alarm goes off, a short press of the wall-mounted button will switch off the buzzer, but the LED will continue to flash. A long press will switch off the buzzer, LED and RLY2, and the system will be re-armed to monitor again for fire.

Note that this is a mains-powered system with no battery back-up. It should always be used in conjunction with traditional battery-powered or battery-backed smoke detectors.

We suggest that this function be activated in all installations since it's unlikely to ever be triggered unless there is a fire.

BCD switches

The four BCD switches that select the various mode, temperature and timer features have internal contacts

DUCTED HEAT TRANSFER CONTROLLER SC ©2025

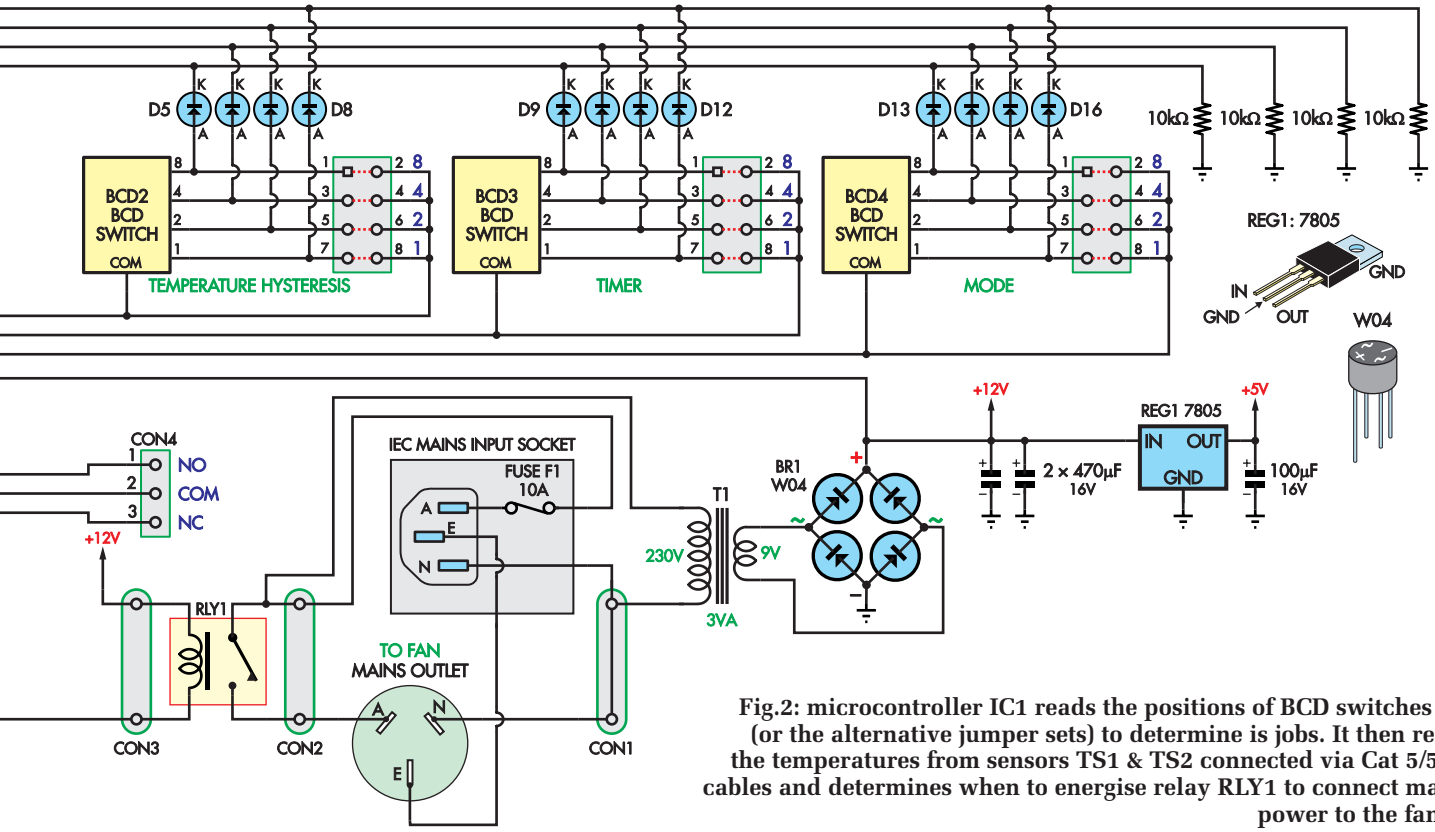

Fig.2: microcontroller IC1 reads the positions of BCD switches 1-4 (or the alternative jumper sets) to determine is jobs. It then reads the temperatures from sensors TS1 & TS2 connected via Cat 5/5E/6 cables and determines when to energise relay RLY1 to connect mains power to the fan(s).

that connect the "1", "2", "4" and "8" terminals to ground in a combination that totals to the switch setting. For example, if the switch is set to the 9 position, the "1" and "8" terminals will be connected to ground but the other two won't.

This allows IC1 to sense 16 possible positions for each switch using four wires (although these switches only have 10 positions).

Rather than the common (C) terminal of each switch being connected to ground, they are connected to a separate pin on microcontroller IC1. This way, the micro can pull them high one at a time, and use the same four lines (RA1, RC5, RA0 & RC4) to read the position of the selected switch.

Isolation diodes D1-D16 are required because, while the other switches can be set to have their common terminals floating while one switch is sensed, those switches could still end up effectively shorting two or more of the sense lines together, depending on their positions.

We need the diodes to ensure the switches don't affect each other during the sensing procedure.

During switch sensing, any open BCD switch will be pulled low to 0V via one of the 10kΩ pull-down resistors.

BCD switches can be expensive, so we have provided an alternative system using a 2×4-pin header with up to four jumpers placed on it to replace each BCD switch.

Fig.3 shows how the jumper settings equate to BCD settings. Since these settings are rarely (if ever) changed, there's little disadvantage in using jumpers on headers instead.

Control Panel

The wall-mounted control panel for the Ducted Heat Transfer Controller comprises switch S1, LED1 and a piezo buzzer.

This is all incorporated in a Clipsal sprung-return rocker switch plate that includes an indicating LED. The piezo buzzer is an addition to the switch plate

#	BCD Switch	Jumpers	#	BCD Switch	Jumpers
0			5		
1			6		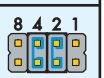
2			7		
3			8		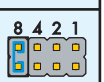
4			9		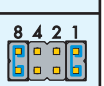

Fig.3: this shows the simple binary codes you need if using jumpers instead of the BCD switches. IC1 also monitors switch S1 and the four selection switches, BCD1 to BCD4. In response to these settings and temperature readings, the microcontroller can sound the piezo buzzer, light LED1 and switch on RLY1 to drive the duct fan. IC1 can also switch on RLY2 if the fire alarm feature is selected with JP1 and is then activated.

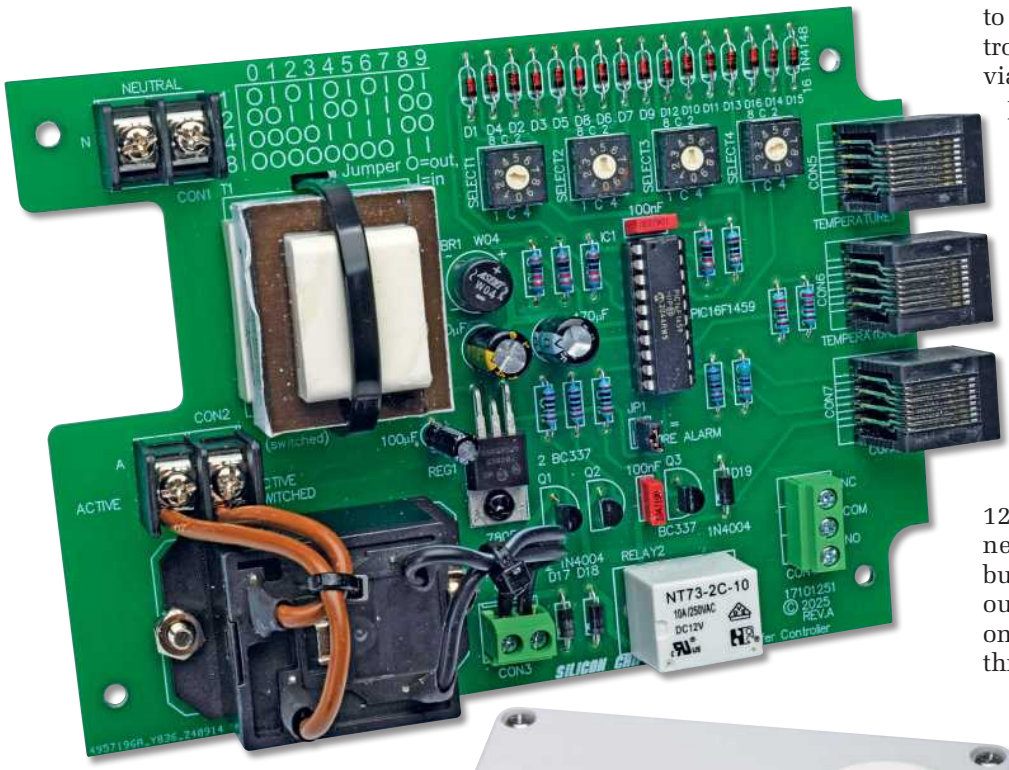

The Ducted Heat Transfer Controller is housed in a polycarbonate IP65 enclosure (upper right photo). An IEC mains cord supplies power and the duct's fan plugs into the power outlet on top. The temperature sensor and control panel connections are made using RJ45 sockets and Cat 5/5E/6 cables. The Controller is easy to build, with only through-hole components used. Care must be taken with the mains voltage wiring, though.

The faceplate (upper left photo) incorporates a momentary rocker switch, piezo buzzer and a white LED that lights the periphery of the switch. The wall plate can be mounted vertically or horizontally – this one is configured for vertical mounting. The 'floppy ears' can be easily removed (they're not needed for normal mounting).

The room temperature sensors are each located in small, ventilated wall enclosures (photo shown at right).

to complete the control panel. The control panel connects to the main board via another Cat 5/5E/6 cable and RJ45 plugs/sockets.

LED1 is driven from the RB6 output of IC1 through a 470Ω resistor to ground. The LED current is around 4.25mA, assuming a voltage drop of 3V across the white LED.

Switch S1 is connected between GND and the RB5 input of IC1, with this input pulled to 5V via a 1kΩ resistor when the switch is open. If the switch is closed, RB5 will be pulled to GND and IC1 can detect that.

The piezo buzzer is powered from 12V using transistor Q3 to switch the negative side to ground. When the buzzer is required to sound, the RC7 output of IC1 is driven high to switch on Q3 by delivering current to its base through a 1kΩ resistor.

Relays RLY1 & RLY2 are switched on via the RC3 and RC6 outputs of IC1, respectively. Both use a 1kΩ base resistor to drive a transistor to power the relay coil. Transistor Q1 is used for RLY1 and Q2 for RLY2. Diode D17, across RLY1's coil, and D18, across RLY2's coil, quench the back-EMF voltage from the coil when these are switched off.

RLY2 is uncommitted and is intended to drive a low-voltage siren for the optional fire alarm function. RLY1 connects the incoming mains Active to the fan socket when the fan should be powered. The output socket's Neutral and Earth pins are permanently wired to the input socket.

Power for the circuit is derived via a mains transformer that produces a 9V AC output. This is rectified by bridge rectifier BR1 and filtered by two 470µF capacitors, giving close to 12V DC. This is used to power the two relays and the piezo buzzer. REG1 is a 5V regulator that drops its 12V input to 5V to supply IC1 and the DS18B20 temperature sensors.

Next month

That's all we have space for this issue. Next issue, we'll cover building the unit and setting it up. PE

Table 2 – smart remote push button/LED/buzzer

Mode	Push button/buzzer	Fan status	Faceplate LED
'0' Manual fan on/off 	**Short press** – beep – on **Longer press** – double beep – off	Runs when fan manually switched on	**Fan off** **Temp difference < set point** LED off **Fan off** **Temp difference > set point** LED flashes momentarily on once every 2s **Fan on** **Temp difference < set point** LED flashes momentarily off once every 2s **Fan on** **Temp difference > set point** LED fully on
'1' Manual fan timed operation 	Quick press or presses = on for set period of operation, e.g. when timer is set for 30m, 1 quick press runs fan for 30m, 5 quick presses sets 'on' period at 150m (2.5h) **Longer press** – double beep – off	Runs for timer period when set	**Fan off** **Temp difference < set point** LED off **Fan off** **Temp difference > set point** LED flashes momentarily on once every 2s **Fan on** **Temp difference < set point** LED flashes momentarily off once every 2s **Fan on** **Temp difference > set point** LED fully on
'2' or more Automatic 	**Short press** – beep – system on **Longer press** – double beep – system switched off	Runs when temperature difference exceeds preset level when system activated	**System disabled** LED flashing **Fan off** LED off **Fan on** LED on
Fire alarm activated (JP1 shorted and fire detected)	Buzzer sounds rapidly and LED flashes rapidly at 5Hz Short press, buzzer sound is off, LED flashes rapidly Long Press, LED and buzzer off and retests for fire	Fan off	N/A

RIGOL DHO924S

Digital Oscilloscope

Review by Tim Blythman

It's been a while since we have reviewed an oscilloscope, and Rigol's new DHO900 series has a slew of modern features with a compact footprint. This review covers these new features and looks at what it's like to use out-of-the-box. We'll also have a brief look at their DM858 digital multimeter.

- 4 channels
- 125MHz or 250MHz bandwidth
- Parallel, UART, I²C, SPI, LIN, CAN protocol analyser
- Sampling rate of 1.25Gsps (shared between active channels)
- 12 bits analog resolution
- Optional 25MHz, 10V peak-to-peak waveform generator
- 7in touchscreen LCD (1024 × 600 pixels)
- Sine, square, ramp, noise, arbitrary waveforms (from CSV file)
- Sample memory of 50 million points
- Period, frequency, rise time, fall time, duty, count, delay, phase and more measurements
- Maths functions, FFT, Bode plot (using the AFG as a source), pass/fail
- USB device, HDMI, LAN, USB host (mouse, flash drive) interfaces
- 15V DC, 3A (65W USB-C PD) PSU

Rigol DHO924S DSO & DM858 meter

After seeing this oscilloscope at the Electronex trade shown in Sydney and being impressed, we obtained a unit to review. The first thing we noticed about the DHO924S is its compact size. For comparison, the older Rigol MSO5354 measures 37 × 20 × 13cm (9.62L), while the DHO924S is only 26 × 16 × 8cm (3.33L). The enclosure is also a striking black colour.

The DM858 digital multimeter appears to use the same case as the DHO924S, so it is a similar size, shape and layout. We also borrowed one of those; our review of it is in a panel later in the article.

The DHO900 series

The DHO900 series comprises four models: DHO914, DHO914S, DHO924 and DHO924S. S-suffixed parts include an arbitrary function (waveform) generator output, while the DHO914 variants offer 125MHz bandwidth against the DHO924's 250MHz bandwidth. The DHO924S has the highest specifications in the series.

Rigol's DHO naming refers to a high-resolution digital oscilloscope; it has a 12-bit ADC (analog-to-digital converter), while many 'scopes we have seen only offer 8-bit resolution. There are four analog inputs, so that's the '4' in the naming scheme.

The DHO924S also has a 16-input logic analyser, but an optional active logic probe is required to use this feature, so we did not test it. We find that four channels are ample for most scenarios, and the analog channels can also be used as inputs to digital features, such as the serial protocol decoder.

There is also a DHO800 series, which offers 12-bit resolution in a mix of two- and four-channel options but lacking the logic analyser option. They have a lower bandwidth and lack some of the other features of the DHO900 'scopes, but are similar in that they have the same compact form factor.

Out of the box experience

The photos show the front and back of the DHO924S. The back panel has two BNC connectors; one of these is for the arbitrary function generator, the other an auxiliary output. By default, the latter emits a pulse when a trigger event occurs. There are also Ethernet, USB-B, HDMI and USB-C connections.

The legs pivot out during use or fold flat for storage. When the legs are extended, the unit leans back about 20°. The legs and base of the unit have chunky rubber feet and the 'scope does not feel like it will tip or slide away while the controls are being operated.

There is an Earth socket for a banana plug, and the brass insets are threaded for M4 VESA mounts with a 100mm spacing. With the DHO924S being so thin, it becomes practical to mount it to an adjustable monitor arm, so it no longer takes up any desk space. Apparently, users are also 3D-printing a variety of adaptors to suit the VESA mounts to adapt to their 'scopes.

The four input BNC connectors are along the bottom of the 'scope next to the ground and test signal clips. The connector for the optional active logic probe is under the display, along with a USB-A socket and the power button.

Compared to the likes of the Rigol MSO5354, there is only one pair of vertical scale knobs, despite there being four channels. The active channel is selected by the same numbered, colour-coded pushbuttons that are used to switch channels off and on. We found this to be intuitive enough and, as it offers more room for the main display, it seems like a reasonable compromise.

The oscilloscope package includes four switchable probes (more on them later), a USB-A to USB-B cable, an Earth lead terminated with banana plugs and a USB-C PD (power delivery) power supply, which indicates that it can deliver 5V, 9V, 12V, 15V or 20V. The oscilloscope uses 15V, according to our USB-C Power Monitor.

The USB-C plug has a pin and latch arrangement that secures it. That means that the power supply cannot be easily used to power other devices, but does not stop another suitable power supply from being used to power the oscilloscope instead.

Being an isolated power supply necessitates the inclusion of the Earth lead, since USB-C does not provide a means for Earthing. In other words, the oscilloscope is floating unless the separate Earth lead is connected.

This is an interesting pitfall that may not be immediately apparent to those accustomed to 'scopes that are normally Earthed via their mains lead. Also, since it isn't directly mains-powered, it is not possible to trigger off the mains waveform without a separate connection.

After powering on the 'scope, it took almost a minute to start up and display the expected screen. Along the way, the message "Android starting" appeared, hinting at the underlying software. The DHO924S we tested ran Android 7.1.2, which dates from 2019.

The LCD has a touch panel, which works as you might expect for an Android device. You can tap on virtual buttons, drag the cursors around and even perform a pinch-to-zoom. There is a "Touch Lock" button to disable the touch panel.

There are several ways to connect to a PC or external monitor. The VESA mounts allow the unit to be mounted on a monitor arm, freeing up bench space. While this is a DHO800 series 'scope, the back of the DHO900 series is identical except for having a black case.

Four PVP2350 350MHz switchable passive probes are included, along with the accessories shown here. The 'scope also includes a power supply & USB cable.

Screen 1: despite the numerous features, it's easy to find most menu options. A good place to explore is the R menu at the lower left corner of the screen. This also gives a good overview of the advanced features.

Screen 2: once we had the 'scope's IP address from this screen, it was easy to connect to the Web Control interface. The Utility menu also includes the setup and self-calibration options.

The 'R' icon at the bottom-left corner of the display opens a menu that includes features beyond what you might expect from a traditional 'scope, such as settings, operating system utilities and the like. Screen 1 shows the contents of this menu.

The front panel looks much like any other 'scope, with the familiar adjustment knobs for vertical and horizontal scale, trigger and RUN/STOP controls, along with other controls to operate the custom features that are seen in modern oscilloscopes. The "Quick" button can be programmed to perform one of several different actions; by default, it saves a screenshot.

The two 'Flex Knobs' near the top do not have a fixed use; their function changes depending on the items selected in various menus. The functions are marked on-screen by small '1' and '2' icons.

The Flex Knobs allow all manner of values to be adjusted instead of being manually entered into a keyboard on the touch panel. This will be handy when values just need to be tweaked by a small amount. These knobs can also adjust the cursors when they are turned on.

Probes

The DHO924S comes standard with four Rigol PVP2350 350MHz passive probes, which are switchable between 10:1 and 1:1 attenuation (as usual, the full bandwidth is only available at 10:1). Each probe includes an assortment of colour-coding rings and a ground spring, as well as the requisite compensation adjustment tool.

The leads are 1.2m long, and their slim cords are light enough to not take up too much space. There is no probe detection (for automatic probe attenuation setting), so these must all be set manually.

Using it

Since we were keen to try out the more modern features of the DHO924S, we hooked up an Ethernet cable. We found the IP address of the 'scope from the Utility menu, as seen in Screen 2. Typing that into a browser's address bar gave us a Web Control Page, allowing us to open a Web Control window that shows the oscilloscope's screen.

Screen 2 was captured on our PC using the Web Control window; it's identical to what appears on the 'scope's

◀ **Screen 3: hooking the 'scope's function generator back to one of its inputs shows off its sensitivity. The 2mV peak-to-peak square wave is the lowest amplitude that it can deliver.**

Screen 4: tapping on each channel's vertical settings shows the full signal path and its associated parameters. You can still see the waveform under the transparent window, allowing the trace to be adjusted with ease.

screen and even gives access to the controls that would otherwise require the touch panel. Any device with a browser and WiFi connection should work. We had no trouble controlling and viewing the 'scope on an Android mobile phone's browser.

This makes it much easier to explore the features of the 'scope, although we think it's a bit of an omission that there aren't controls for the various hardware buttons and knobs. Still, most settings can be set via menu panels.

The HDMI interface just works. We plugged in a HDMI cable, connected it to a monitor and the 'scope's display appeared on the screen without having to change any settings. The output appears to be identical to that on the LCD, scaled up to use the entire viewable area of the monitor.

We'll delve further into the various interfaces a bit later, including some software that Rigol offers. The relevant downloads can be found at www.rigolna.com/download/

There is a self-calibration mode that is recommended to be performed after the 'scope has warmed up to operating temperature (after about 30 minutes). The self-calibration process took about 24 minutes to run on the unit we were testing.

The sampling rate is shared between the four channels, since there is only one analog-to-digital converter IC. Using two channels halves the available sampling rate, while using three or four channels will reduce it to a quarter. So the maximum sampling rate can only be achieved if only one channel is in use. Similarly, the sample memory is also shared between the channels in use.

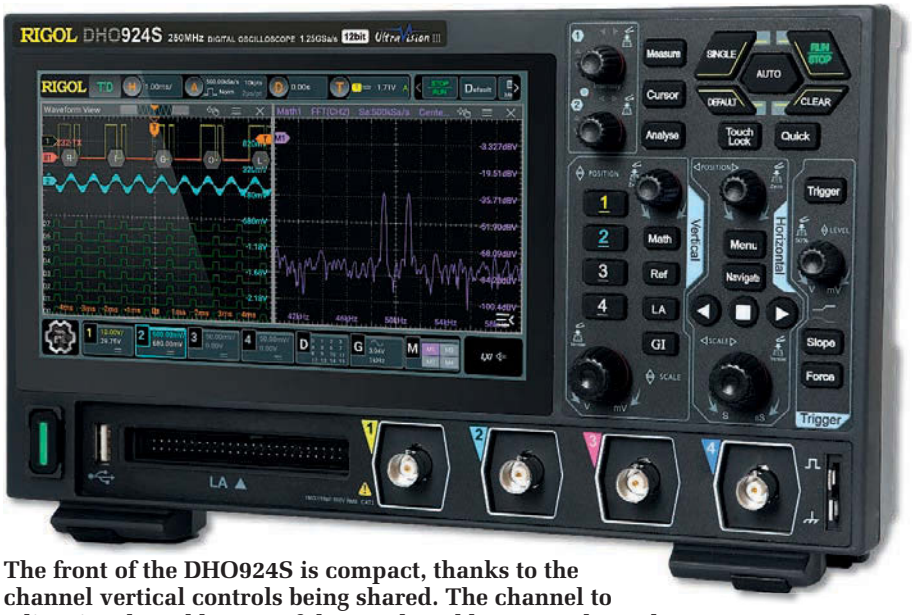

The front of the DHO924S is compact, thanks to the channel vertical controls being shared. The channel to adjust is selected by one of the numbered buttons. Above these are the Flex Knobs, which adjust values depending on the current sub-menu.

Noise performance

Using the inbuilt arbitrary function generator, we looped a 2mV peak-to-peak 1kHz square wave signal back into the 'scope with no attenuation and with the bandwidth limited to 20MHz. The result is seen in Screen 3. Some of this noise will be from the function generator, but clearly, the DHO924S has no trouble resolving signals below 1mV, which is pretty impressive.

Inputs

Opening the menu options for the inputs reveals the signal flow diagram seen in Screen 4. Some of the options can be adjusted by tapping on the diagram itself. You can use the Flex Knobs, as indicated by the yellow hexagons.

Screen 5: measurements are added to the Result panel at right. This panel can scroll up and down, so more than five results can be added. The options relating to the horizontal (time) axis are shown here.

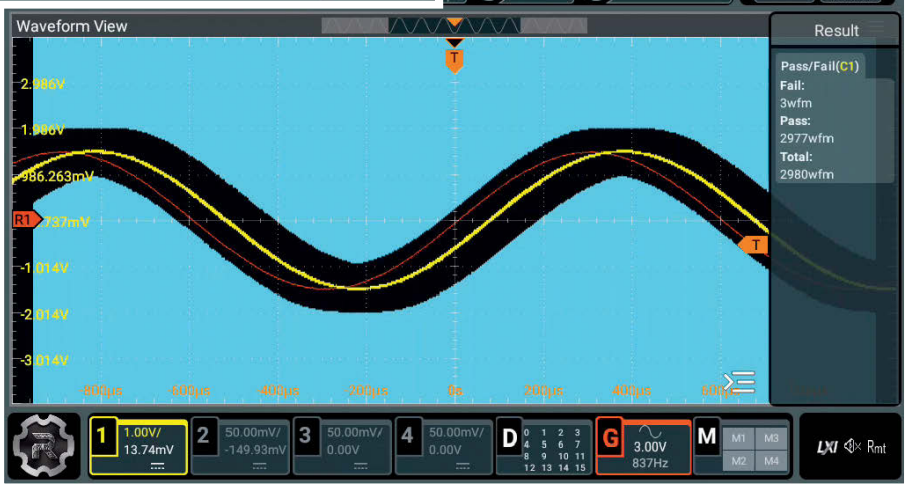

Screen 6: this display has both reference waveform (orange) and pass/fail mask (blue) active. The auxiliary output at the rear of the 'scope can be set to produce a pass/fail signal.

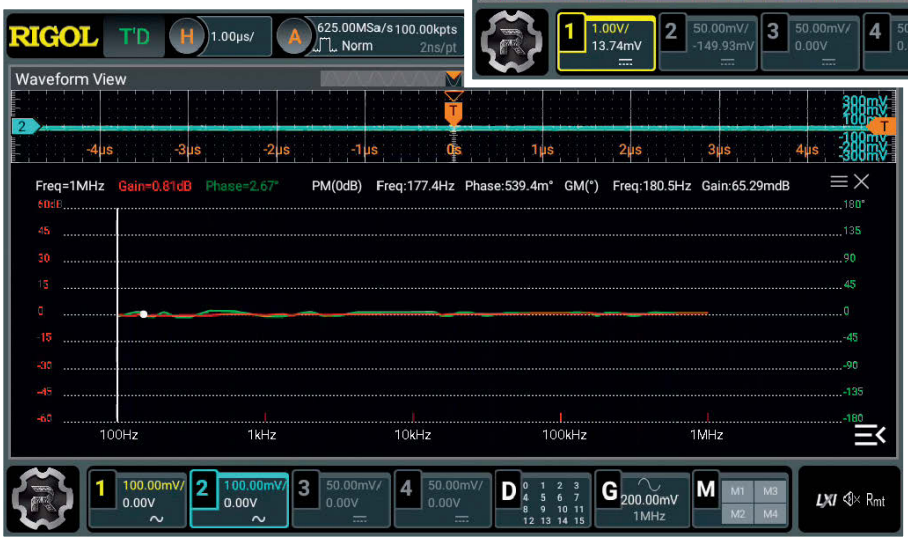

Screen 7: a Bode plot expands to take up most of the available screen space. Most of the windows are movable and adjustable, so you can customise the viewport to suit your preferences.

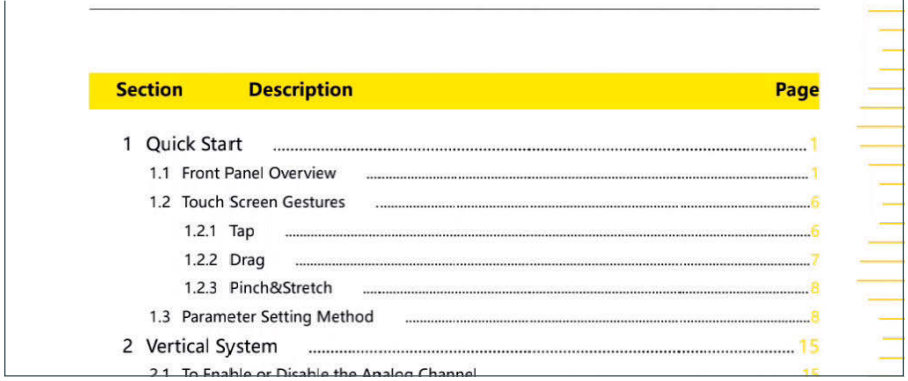

Screen 8: the 'scope's help system includes a PDF copy of the manual that can be viewed on the 7-inch display.

For numerical values, a pop-up keyboard can also be used to enter values.

The main parameters are also displayed on the channel widget near the bottom of the screen, and you can see the appropriate parameters for the arbitrary function generator displayed on the G channel too.

Maths and measurements

Handy features on most modern 'scopes include various mathematical functions and trace measurements; the DHO924S is no exception. The measurement menu can be accessed from a fixed button on the right-hand side or via the 'Measure' button at the top-right of the screen, on the touch panel or even via the main 'R' menu button.

The area at top-right actually hides several different buttons that can be accessed by swiping left or right. Fortunately, most menu items can be accessed in different ways, so you can choose whatever option is most intuitive.

There are 41 different measurements available; Screen 5 shows some of these active, as well as the horizontal (time-based) measurements that are available. The vertical options include voltage-based (peak-to-peak, RMS etc) measurements, while the remainder are time and phase delay measurements between two input channels.

The Setting option allows thresholds to be set. These default to 10%, 50% and 90% levels of the waveform, which worked quite well for us during our tests. You can also view various statistics, such as average, minimum or maximum of the measured parameters, or view a histogram of the measurements as they are gathered.

Four 'Math' channels are available, including operations such as summing or differencing two channels. Single-channel operations such as logarithm, exponent, derivative, integral and square root are available. FFT (fast Fourier transform) or spectrum analysis is also possible.

Functions

The menu shown in Screen 1 gives an idea of the DHO924S's built-in functions. When measurement or maths windows are open, they can be dragged around and reorganised (something we're not used to on a 'scope).

A Reference waveform can be captured ("Ref") from an active channel to be visually compared with the changing input. A more rigorous signal check can be performed using the Pass/Fail function. This requires a mask against which the active channel is compared; the results of the Pass/Fail can optionally be fed to the auxiliary BNC connector on the rear of the scope.

The mask can be loaded from a file or can be easily created by applying margins (in time and voltage) against a sampled input channel. A typical example of such a mask is the so-called eye diagram used to verify high-speed signals, such as USB or HDMI. Screen 6 shows the Reference and Pass/Fail functions.

The protocol decoding feature supports Parallel, UART, I²C, SPI, LIN and CAN signals, although you would probably need the optional active logic probe to do much with a parallel bus. There are many parameters available to adjust, including polarity, parity and bit order, although the defaults look to be sensible for commonly used configurations.

There is a button to swap SDA and SCL in I²C mode, so it's reasonable to just hook the 'scope up without worrying too much about which wire is which. It can decode up to four separate buses at once, which is more than sufficient for most applications.

The Bode plot function is only available on models with an arbitrary function generator, since the generator is used to provide the input waveform. The Bode plot window expands to take up most of the display, as seen in Screen 7.

The Auto function is the same automatic configuration utility that is found on other 'scopes to quickly set up the timebase, voltage scale and trigger selection based on the active signal. It can also be accessed from the hardware button in the top right corner of the 'scope.

The Display settings can select trace persistence and change other parameters as trace and grid intensity, as well as window transparency. On the bottom row are functions related to the operating system functions of the unit. The Help feature is actually an on-screen viewable PDF version of the manual, which can be seen in Screen 8.

Other connections

The front USB-A socket can be used for either a mouse or flash drive (a USB hub allows both to be connected

PoLabs
For more information or software download please visit
www.poscope.com/epe
PoKeys
Connect, Control

- USB
- Ethernet
- Web server
- Modbus
- CNC (Mach3/4)
- IO
- PWM
- Encoders
- LCD
- Analog inputs
- Compact PLC
Stepper motor drivers
Drive
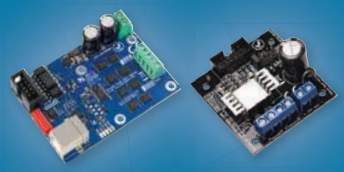
- up to 256 microsteps
- 50 V / 6 A
- USB configuration
- Isolated
- up to 32 microsteps
- 30 V / 2.5 A
PoScope Mega1+
PoScope Mega50
Measure

- up to 50MS/s
- resolution up to 12bit
- Lowest power consumption
- Smallest and lightest
- 7 in 1: Oscilloscope, FFT, X/Y, Recorder, Logic Analyzer, Protocol decoder, Signal generator

simultaneously). The mouse is used as you might expect, to interact with items on the screen. We think that support for a keyboard or numeric keypad might be handy for parameter entry, since this can sometimes be awkward to do with an on-screen keyboard.

A USB flash drive can be used to transfer screen captures or waveforms for the arbitrary function generator. It can also be used to upgrade the firmware. The 'scope has internal storage and can connect to an SMB (network) file server via Ethernet.

We often use screen captures or 'scope grabs for articles in the magazine, so we thought we would use a USB flash drive to transfer the necessary files for this article. But it turns out that the web control and network interface mean that is unnecessary, since we can download captures directly to a PC.

The Web Control's Print Screen tab has buttons to take a screenshot or record a video; the image or video is displayed in the browser window and can be simply downloaded onto the PC from there. The older MSO5354 also has a web control interface, but we found that quite slow; in comparison, the DHO924S feels much snappier.

The USB-B socket on the rear of the 'scope can be used to control the DHO924S, but we found that the Ethernet connection was more useful, since the PC does not need to be near the 'scope as is required for a USB connection.

Other notes

There is a sleep mode which can be used instead of a full shut-down. This has the advantage that the 'scope only takes about 20s to be ready from sleep, and also retains its last operating state. The main downside is that the power supply must remain active to retain this state.

Using our USB-C Power Monitor, we recorded a peak of 2.6A at 15V, consistent with the 3A maximum recommended in the user guide. During sleep, we recorded a draw of around 230mA at 15V, while a full shutdown reduces this to about 1mA.

The web control view would occasionally reset, changing the window size and stealing the focus from another window if it did not already have focus. If the network connection is lost, a message is displayed in Chinese. The same message is shown if a second browser window attempts to connect to the web control.

Software

SCPI (Standard Commands for Programmable Instruments) is a standard for the control of test equipment and instruments. The DHO924S presents an SCPI interface on Ethernet port 5555, and it can also be accessed via USB.

Rigol provides a Programming Guide which outlines the SCPI commands specific to the DHO900 series if you wish to control the DHO924S this way.

Rigol also provides software that can connect to its 'scopes and other hardware. The Ultra Sigma software can connect to the DHO924S via Ethernet or USB, and has an SCPI Panel Control, which can send commands and read data using the SCPI interface. Screen 9 shows the main window for Ultra Sigma with both the DHO924S and DM858 connected.

Ultra Sigma appears to be only available for Windows, although the web control should work on a browser under most operating systems.

We have heard, but it has not been confirmed, that Rigol will release updated PC software at some later date. However, we will not hold our breath as it has been a while since we heard that.

Overall impressions

With the numerous menu options, it was easy enough to find a specific function and everything feels intuitive. The 'scope feels like it has all the features we might need and probably a few we don't realise we need yet.

All the controls of the DHO924S feel quite snappy and responsive, whether using the knobs and buttons, the touch panel or the web control interface, although the waveforms will freeze while dragging.

We didn't often touch the wrong item on the touch panel, since most objects are quite large, but it happened at times, and felt slightly fiddly. Using a mouse was much more precise, so that is a good option. We mostly used the web controls for much the same reason.

The noise level, bandwidth and number of channels makes this oscilloscope suitable for a wide range of jobs, for which an expensive high-end 'scope would have been required in the not-too-distant past.

Conclusion

If we were looking for a 'scope right now, the DHO924S would definitely be on the shortlist. The web control and Ethernet interfaces make it very easy to capture screen grabs and other waveforms and analyses. It also makes it easy to control the many functions of the 'scope.

The unit is compact and light. It's responsive and intuitive to use, and most of the specifications easily exceed the 'scopes that we currently use.

The DHO924S is available for £744 including VAT: https://www.rigol-uk.co.uk/product/rigol-dho924s-4ch-12bit-250mhz-1-25gsa-s-digital-oscilloscope-25mhz-arb-gen/ PE

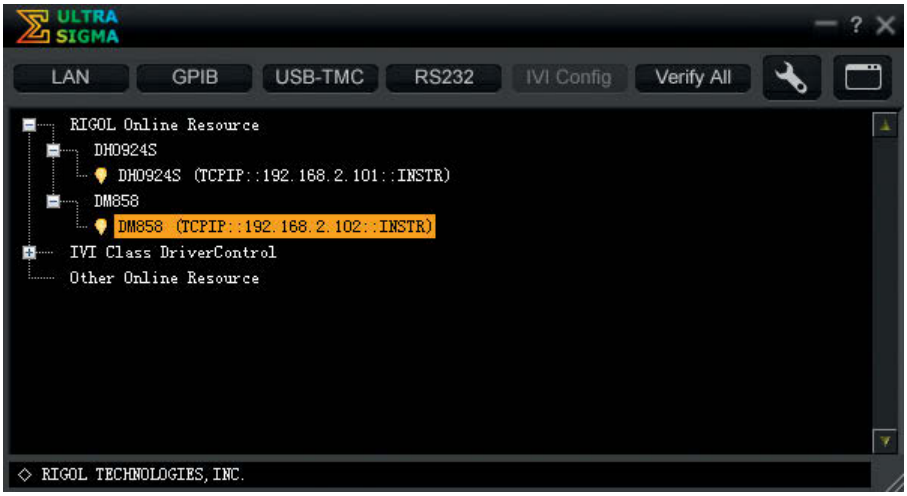

Screen 9: the Ultra Sigma program can interface with many Rigol instruments. Here, both the DHO924S 'scope and DM858 benchtop multimeter are connected. Ultra Sigma includes an SCPI control panel.

Rigol DHO924S DSO & DM858 meter

DM858 Digital Multimeter Review

We haven't had much need for a benchtop multimeter, with the handheld variants being sufficient for most purposes. A benchtop multimeter falls between a handheld multimeter and an oscilloscope, including a 5.5 digit display capability and features like those you might normally see on an oscilloscope.

The DM858 comes with a power supply and a pair of CAT II multimeter leads, as well as a pair of alligator clip adaptors that screw onto the ends of the probes. Like the DHO924S, the power supply is a 65W USB-C PD (power delivery) PSU. The DM858 only requires 12V at up to 10W, with our USB-C Power Monitor indicating a typical draw of just 7W.

The back panel connectors are almost the same as the DHO900 series 'scopes, with only the HDMI socket missing. It has a USB-B socket for connection to a computer, and an RJ45 jack for Ethernet. One BNC socket is for an external trigger input, while the other can output a pulse after each measurement. Since it is much the same case, the same 100mm VESA mounts are present.

Using the DM858 is easy for anyone who has used a multimeter. The extra banana sockets allow four-wire (Kelvin) resistance measurements.

Apart from the HDMI output, most of the user interfaces are the same as the DHO924S; the Ethernet and USB connections on the back of the unit can be used for remote access. Web Control for the DM858 works similarly, as does access to the SCPI interface over Ethernet.

The front USB-A socket supports a USB flash drive or a mouse, and it has five banana sockets on the front panel, along with the fuse for high-current measurements.

Three of these sockets are used as you might expect, with one common input used in conjunction with another input for high-range current measurements. The third input is used for all other measurements, such as voltage, resistance, capacitance and so forth. The other two connections can be used for Kelvin (four-wire) resistance measurements. The Kelvin technique is often used to measure low resistance values, since it eliminates contact and lead resistance that might interfere with the measured resistance.

In use

The DM858 takes about a minute to boot up and it then shows a DC voltage reading. The default, slow update rate is easy to follow. There are two faster update rates available.

There are buttons for typical multimeter measurements, such as voltage, current (DC and AC), resistance, continuity, diode, capacitance and frequency; standard operations are fairly obvious. Some features, including four-wire resistance measurement and diode mode, are accessed by pressing the Shift key.

The overall interface is very similar in feel to the DHO924S, with a 7-inch touch panel offering menu items above and below the main display. An 'R' menu in the bottom-left corner offers a range of functions that duplicate some of the controls, besides allowing access to the system controls and settings.

Other features

Despite being labelled a multimeter, the DM858 can display a slow-moving trace, but it's more like a chart recorder than an oscilloscope. The fastest update rate is around 10Hz. There are simple 'Math' functions, such as applying an offset to a reading, as well as statistical results such as minimum, maximum, average and standard deviation. Voltages can be converted to dB values.

The DM858 can also interface to sensors such as thermocouples and thermistors, with several inbuilt probe types and presets being provided. An adjustment for cold junction temperature can be added.

Other custom sensors can be monitored by supplying a list of measured value (resistance, voltage, current etc) and display value (such as temperature) pairs. This makes it possible to set the meter up as a custom display to suit just about any type of sensor.

Summary

We found the DM858 easy to use. It offers numerous handy features above those of most multimeters. We would make good use of the Web Control interface to allow remote viewing and operation. The DM858 digital multimeter can be purchased from https://www.rigol-uk.co.uk/product/rigol-dm858-5%C2%BD-digit-digital-bench-multimeter/ for £458.40 including VAT.

The DM858 is small enough that it can be mounted using a VESA mount. It also has an Ethernet connection on the rear, as shown in the photo above.

Net Work

Alan Winstanley

Our occasional column delves into popular 21V rechargeable tools and batteries, brings updates on the forthcoming 'Temu Tax' and ruminates about buying an electric vehicle (EV).

In recent years, rechargeable power tools and equipment have become widespread. Popular brands such as Ryobi offer consumers some 200 products for the home or garden, all powered by the same "One+" 18V battery pack they invented 30 years ago.

The idea of offering swappable batteries came not a moment too soon, as many DIY tools of that era were rendered useless once their NiCd or NiMH batteries lost capacity, or they just fizzled out altogether and had to be scrapped.

Professionals use rugged and expensive power tool brands like Makita or DeWalt that withstand demanding trade use. For example, I saw a pair of high-end 8Ah DeWalt batteries cost nearly £250 alone, and smaller 4Ah types were still £50 each.

Ryobi-brand 18V One+ batteries for consumers aren't particularly cheap either, and thus compatible clone-types have become popular. I bagged two such 8Ah batteries for £25 on offer from a Temu Local Warehouse seller.

In pursuit of my horticultural interests, I used them with a Ryobi USB converter to create a hefty 5V standalone supply to power air circulator fans in my greenhouse during this summer's heatwave. The twin-blade fans also came from Temu, a snip at under £7 each.

Strangely, the Ryobi USB converter (type R18-USB-0) doesn't indicate when it's powered up, so for fun, I added a super-bright wide-angle multi-colour LED, as described in my blog at **https://pemag.au/link/accq**

There are many budget-priced alternatives to these rechargeable tools and batteries. A quick search of popular Chinese tech websites soon reveals a large choice of "21V" power tools and Li-ion batteries.

In my *Net Work* column in the December 2025 issue, I highlighted a typical battery-powered pressure washer that utilises these 21V power packs. It's left to individuals to decide whether to spend £35-45 locally on these if they could import directly from China for £12-15 while they still can.

In my case, the battery pressure washer was followed by a perfectly satisfactory mini chainsaw kit with two batteries and an air blower – all at affordable prices; little 'wins' that are welcome in today's straitened economy. But as I describe later, the honeymoon may soon be over.

Rounding off my unbranded 21V battery toolkit (for now, anyway!) were some small LED spotlights that even YouTuber BigClive-DotCom thought were decent enough, plus a small USB-C power bank adaptor. An unusual surface-mounting adaptor plate allows a battery to securely 'click' onto electronic enclosures or other projects, although an inline fuse is strongly recommended.

Genuine Makita batteries have an LED 'fuel gauge' and AliExpress supplied a small four-LED voltage monitor PCB that fits certain 21V battery packs directly, to provide a push-button LED battery checker.

Some batteries have a rounded rectangle 'placeholder' moulded on the plastic housing for this; the PCB is a perfect drop-in fit inside (see below). It's a pity that the housing is blanked off where the front 'control panel' would fit, though. Otherwise, the separate plastic actuator for the switch could simply slide in, in front of the PCB, and then the stick-on bezel would finish off the job nicely.

Retrofitting the fuel gauge onto a 21V battery will be a job for a rainy day!

Faking it

China's dominance of Li-ion battery manufacturing has been instrumental in enabling these power tools and equipment to be mass-produced. The 21V battery pack design has been heavily cloned. Most of them are unbranded and untraceable, and they carry little, if any, certification or labelling. Even so, it doesn't mean that they are not fit for purpose, but there's always 1% of remaining doubt.

It turns out that the battery design

A USB adaptor on a 'clone' battery, ideal for powering USB-C fans in the author's greenhouse this summer.

This plate enables a 21V Li-ion battery to act as a clip-on DC power source. An inline fuse is essential.

An LED 'fuel gauge' PCB sourced from AliExpress intended for Makita G-series style Li-ion battery packs.

These fans, powered by a tool battery, came in handy as temperatures under glass hit nearly 50°C!

was ‘inspired’ by Makita’s budget price, low-end G-series power tools, not to be confused with their mid-range performance LXT tools and batteries used by professionals.

The plasticky, plug-in mains battery chargers are heavily cloned too and, judging by their weight, contain mostly fresh air. Although my expectations were low, they do the job; still, I tend to keep half an eye on them when in use. A Kowsi DC-28 DC plug-in monitor helps display progress (as described in the November 2025 column, p56).

Mechanically, I feel these 21V batteries (dubbed “Fakita” by Big Clive) are adequate for their intended domestic use, in keeping with their low-end aspirations. Two slots on the battery contain the positive and negative poles, while sometimes a smaller third slot, usually unpopulated, would be for a thermistor battery temperature monitoring connection found on authentic Makita tools.

There’s no escaping the fact that these budget, low-power battery packs are churned out by the million, built using commodity 18650-size Li-ion cells. Their marketing has one major problem though: the batteries are not 21V, but 18V!

The Kowsi DC-28 voltage and current monitor is ideal for keeping an eye on a battery’s charging status.

Inside a 21V Li-ion battery. It holds five commodity 18650 cells. The main PCB is behind.

These mass-produced “21V” battery packs contain five 3.7V lithium-ion cells. This produces an 18V nominal working voltage. The stickered “21V” is simply the battery’s peak off-load voltage. The “21V” labelling is deceitful marketing that makes 21V tools appear more powerful than mere 18V ones – a con trick that will sway ordinary non-technical consumers.

I inspected the innards of a typical 21V 2Ah battery (see photos). The battery pack pops out and the spot-welding between the five 18650 Li-ion cells was of consistent quality. To my eye, the electronics assembly was satisfactory. Mounting slots inside the case would hold the LED fuel gauge that I mentioned before.

As these designs are widely cloned, expect the quality to differ between sources, but the ones I checked seemed adequate. However, the bulky-looking 18V “6.0Ah” battery was a fake. The 10-cell holder inside contained five cells, wired identically to the 2Ah one shown.

These power tools and accessories are widely available through Amazon, eBay, AliExpress and Temu – also look out for Temu Local Warehouse offers which ship from within the UK, a shopping trend that is on the increase.

Importantly, Li-ion cells must be disposed of properly and never thrown into the wheelie bin to avoid the risk of causing a fire in the garbage truck or a conflagration back at the depot.

Spot the difference: an honest 18V Li-ion battery or a common “21V” type. They are the same thing.

Coming soon: the Temu Tax

In the last *Net Work* column, I mentioned how the de minimis rules governing goods imported into the UK were set to be abolished. Currently, orders up to the value of £135 are exempt from customs duty, a fact that is exploited by Chinese e-commerce vendors such as Shein and Temu.

These are known as LVIs, or Low-Value Imports. Some analysts estimate that Temu alone ships at least 100,000 LVIs into Britain every single day. The UK is their No. 2 market after the USA, and in Europe, Britain is the runner-up to Germany, which is estimated to import 400,000 Temu and Shein consignments daily.

A 2025 estimate reckoned 600 million LVIs arrived here in one year alone, and Temu shipped over two million parcels worldwide daily. Parcels that come overland face a tortuous road trip from China to Kazakhstan, through Poland, Germany and Holland before arriving in Britain.

The same issue of fending off cheap imports has struck the European Union, which has faced a

For Free Magazines
sastatus.com

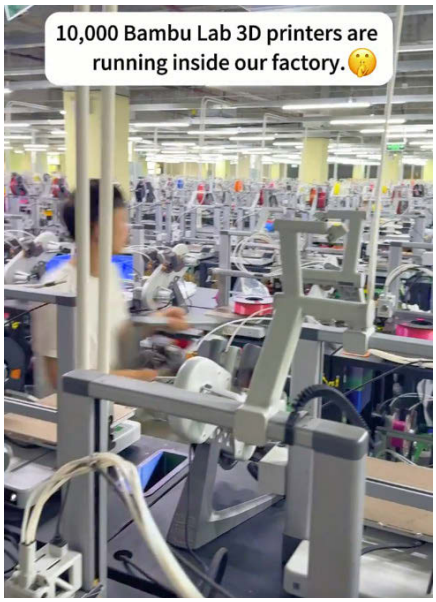

One Chinese plastics manufacturer boasts of having 10,000 3D printers under one roof!

blizzard of low-cost e-commerce imports from China. Just as in the UK, poor-quality Chinese goods that fail to meet safety standards are an ongoing concern. Deliberately under-declaring their value in order to evade customs duty is another problem.

The EU €150 de minimis customs duty exemption was therefore abolished in July 2026, and until the EU makes up its mind about the way ahead, a 'temporary' flat rate of €3 per line item (not per parcel) has been applied, although many sellers are actually applying this per item at the moment.

The EU levies this on e-commerce orders delivered to EU consumers from 'third countries' (EU-speak for everywhere else but the EU, such as China or the UK; note there is no such thing as a 'second country').

This levy is theoretically based not on quantity but on the item's tariff number or unique commodity code. Under the new rule, if consumers import into the EU a pair of jeans, some shoes and a soldering iron, for instance, then three duties are applied, totalling €9, as three different commodity codes are involved. But a solitary soldering iron would cost just €3.

Needless to say, electronics hobbyists have been outraged by duty imposition on cheap electronic parts sourced from the likes of AliExpress, especially given that many of these products are made in China and not really available elsewhere.

Britain's now-deposed Chancellor, Rachel Reeves, left a ticking time bomb in her Budget announcement when she stated that the UK's de minimis rules would be scrapped. In Britain, there are moves to bring forward the new import duty rules to raise badly needed revenue and combat the onslaught from Chinese e-commerce sites.

It would also go some way towards aligning LVI trade with freight that arrives through more traditional shipping channels.

All indications are that the rule will be implemented by October 2028 at the very latest. My best guess is that a levy of £3 per commodity will probably be charged. The duty is collected by the seller, and don't forget that VAT is charged on top.

Once consumers feel the burden of the extra duty (plus VAT) levied on their small orders, the new rule will probably dampen trade in low-value imports. We can expect many more 'Local Warehouse' vendors to spring up, UK-based sellers that avoid consumers being hit by the consolidated import duty. In the meantime, buy now, while you still can!

Plastic surgery

One of the joys of having a lifetime career enriched by industry is that you finish up knowing a little about an awful lot of things, and an awful lot about one or two things! Among other topics, as an industry buyer of several million plastic mouldings per year, I became acquainted with much of the technology of plastics manufacturing.

At that time, a neighbouring factory moulded our products for us, and I would often find myself admiring some of their fascinating injection-moulding equipment. The firm in question acquired another press (the moulding machine) from an auction, and I was soon helping them to set up the thermocouples and temperature controllers.

The factory owner never did grasp my explanation of Ohm's law, which he invariably called "Home's Law"!

The steel mould itself is simply called the 'tool', and the firm's toolmaker manufactured the complex moulds from blueprints. So-called multi-impression tools produce many identical items at once; 24 or 48 aerosol cups, for example.

Fun fact: if you see a solitary number stamped on a plastic product, such as 2 or 26, that's probably the impression number. If some mouldings are somehow faulty, the operator knows that No.2 or No.26 needs attention.

Machines known as spark eroders (or electrical discharge machines) nibbled away at steel blocks in an oil bath to produce the complex shape of the plastic mould. Actually, the factory boss had a radio playing nearby, and a crackle of interference indicated the spark eroder was working properly. Later, CADCAM machines streamlined the tool-making process.

It was at this time that the factory owner complained that he was starting to lose business. Even his regular customers were starting to source plastics from China, where tool-making and production were much cheaper, and it was impossible to compete. Gradually, his plastics moulding business withered away due to intense competition from Asia.

We then brought our plastics moulding production in-house. When I visited recently, I learned of the crippling effects of the UK's ruinous energy pricing was having on industry.

Over in China today, plastics moulding suppliers are everywhere, and CNC machines produce the tools with alarming efficiency. But there is still the cost of tool-making to bear (say £2500-£12,500). Even at Chinese prices, low-cost production runs of smaller plastic parts are unaffordable. Which brings me neatly to the topic of 3D printing.

A decade ago, in the December 2015 issue of *EPE*, the writer reviewed the Dremel 3D Idea Printer, a re-badged machine from Bosch-owned Dremel. This fully enclosed, microwave-oven sized single-filament PLA printer was aimed at the emerging hobby of 3D printing, though the subject had been touched on even earlier in 2009, when we reviewed the Rapman 3D printer.

3D printing would become a fascinating and rewarding hobby in itself, with much of the challenge going into designing 3D objects on-screen before sending them to the printer. A much wider range of filaments is now available, and of course, 3D printing is ideal for low-volume or one-off projects.

Chinese manufacturing has since taken 3D printing to a whole new level, with some 3D tech firms employing many thousands of machines all under one roof. 3D printing of plastic or metal parts

has replaced expensive injection-moulding or tooling with computer design files instead.

One factory shown on TikTok claims to have 10,000 Bambu 3D printers that can mass-produce parts for toys and such with alarming ease. E-commerce sites such as Temu flog numerous handy little plastic aids and doo-dahs that turn out to be 3D-printed at very little cost. The quality is adequate, although some turn out to be poorly designed and break surprisingly easily.

Chinese manufacturers such as Shanghai Yunyan (**ymolding.com**) offer fast turnaround of 3D printed prototypes, moulding, CNC machining, sheet metal parts and much more. I'm not sure what my plastics injection moulder factory owner would have made of it all when he pondered the meaning of "Home's Law" while we wired up the controls of his secondhand machine, bought cheap at an auction.

Getting ready to switch?

A neighbour recently asked me what I thought about buying an electric vehicle. I reckon that the roads will look very different in 20 years' time, with EVs dominating the market. By then, the internal combustion engine may nearly be consigned to history, notwithstanding the prospect of driving hydrogen-fuelled cars.

Still, as far as buying current EVs was concerned, I feared we were currently going too far, too quickly.

One Chinese industry consultant described how, when EV maker Tesla opened its first gigafactory in China, they "accidentally taught the Chinese how to make electric vehicles". *[That was no accident – China's longstanding policy on vehicle manufacturing deliberately used foreign manufacturers and joint ventures to develop domestic expertise – Editor]*

Regular *Net Work* readers will have seen how the EV market in Britain and elsewhere has turned the motor car sector on its head; arguably, it's almost on its knees.

After a faltering start, as well as meeting lots of resistance from die-hard petrol and diesel owners, the signs are that EV brands are now gaining traction. During one local shopping trip, I saw no less than five Jaecoo EV cars, Range-Rover size vehicles that ooze showroom appeal and high-end technology. And this was in the provinces of rural England, not in the suburbs or cities.

The Ora Ballet Cat EV by Great Wall Motor is legally if not aesthetically distinct from the VW Beetle.

The Chinese commentator likened EVs to 'smartphones on wheels'. I then recalled how motoring correspondents commented years ago that there wasn't much room left to develop the mechanicals of a fossil-fuelled car much further. Any technological advancement would likely come from adopting electronic controls and peripherals instead.

This prospect has now come to pass, as the latest EVs overflow with technology but have few moving parts. Smartphones on wheels, indeed!

When the classic English MG motor brand fell under Chinese ownership, it got a head start in the British EV market. Only three years ago, back in the May 2023 issue, I mentioned a barely known EV upstart (in Britain, anyway) called "Build Your Dreams". Their first entry into the UK EV market would be the striking BYD Atto 3.

In no time at all, BYD has built the largest factory in the world, together with a fleet of super-size car transporter ships that deliver 9000 EVs at a time, all around the world. I've now lost count of the number of Chinese EV brands that have stormed into the UK market since then. Even my regular Honda dealer now sells BYD alongside the Japanese marque.

BYD has the advantage of also making its own 'blade' batteries, and new technology sodium-ion batteries are promised. I understand that BYD was making its own silicon chips nearly 25 years ago, long before it built its first car.

BYD announced in May its first four-nanometre smart driving chip, destined to form the heart of autonomous EV control. The Xuanji A3 chip is designed entirely in-house and is a sign of how BYD confidently sees the future. There are more details on this in the video at **https://youtu.be/HXPPBVF2yPE**

There's no disputing the technological prowess of Chinese EV makers. We now face the risk of the entire motor sector being engulfed by these vehicles. Visitors to showrooms find the tech very alluring, even if a vehicle's sloppy dynamics and handling are decidedly un-European-like.

Some owners voice concerns about persistent faults, insurance costs, rust, repairs and spares, and the effects of depreciation have yet to hit home.

Plenty of Jaecoo owners are enjoying driving what some dub the 'Temu Range Rover' though, as my recent supermarket dash demonstrated. I've spotted several BYD cars on the road as well.

There are still plenty of bumps in the road ahead. In November 2023, I showed the GWM (Great Wall Motor) Ora "Funky Cat" compact EV, whose cutesy looks would, they hoped, find favour in the British market. In April this year, GWM pulled the car from the market, citing very poor sales. The importer of Skywell cars also ceased trading this year, leaving the brand's distribution in limbo.

Ironically, GWM continues to sell a retro-styled VW Beetle lookalike in China called the Ora Ballet Cat. It was the British Army that revived car production in post-war Germany, with the Volkswagen Beetle being manufactured initially for Army use.

Such is the perilous state of today's car market that Volkswagen is reportedly looking at slashing up to 100,000 jobs, hiving off some brands, closing factories or diversifying into military production instead.

As I said to my neighbour over coffee, if you twisted my arm, I might consider the Renault 5 or Renault 4, but thinking about it is as far as I've got.

You can contact the author at **alan@epemag.net** PE

Part 2 by Richard palmer

HOME ASSISTANT

WITH A RASPBERRY PI

Having set up a custom Home Automation (HA) system, we now look at advanced features like cameras, custom dashboards, IR remote control, notifications and remote access on a mobile phone, tablet or PC.

Last month, our article on Home Automation explained how to set up a Raspberry Pi to act as a Home Automation hub. We also presented a project article to build a Satellite, allowing the connection of all sorts of low-cost sensors, relays, displays and more.

We mentioned some cameras that can be connected to this HA system but didn't give instructions on doing so. We will now explain how to connect a few different types of camera.

After that, we will create a custom dashboard, a temperature history chart and an adjustable thermostat. We will then look at switching devices on and off with infrared remote control, set up remote access from a smartphone or tablet using a VPN and create an intruder alert with a passive infrared (PIR) sensor detects motion.

Finally, we will explain how to back up your HA system in case something goes wrong.

Setting up cameras

Three of the options for adding a camera to HomeAssistant are explored below: a USB camera connected directly to the HA hub; a $10 ESP-CAM converted into an ESPHome webcam, and an ONVIF compliant IP camera with pan and tilt functions. Each has advantages and disadvantages.

The UVC-compliant USB camera was the easiest to configure, but must be plugged directly into the HA hub. It provided good images in daylight.

The ESP-CAM has WiFi and was the least expensive and straightforward to configure. Daylight image quality was limited, but good enough for many purposes.

The ONVIF IP camera was the most complex to set up, requiring a smartphone app to connect it to the network. Two of the four IP cameras we obtained were not ONVIF-compliant, despite their specifications saying they were.

Once connected to the network using the mobile phone app, the two ONVIF-compliant cameras were automatically set up by HA. These cameras produced the best pictures overall, particularly in poor light or darkness. Both had pan and tilt capabilities that could be accessed from HA.

USB camera setup

Plug the USB (UVC) camera into a spare port on the Raspberry Pi. There is no need to switch the hub off to do this.

Click on your name at the bottom of the left menu bar. Enable the "Advanced mode" slider in the first block of options. Go to Settings, then System, then Hardware and click on the ALL HARDWARE link. Type "video" into the search box.

Click on the down arrow next to the first entry that resembles "video0". Copy the "Device path:" value. It will be something like "/dev/video0".

Exit the menu and go to Settings, then Add-ons. Install the "File editor" add-on by going to the ADD-ON STORE and searching for "file". Once it has finished installing, enable the "Watchdog" and "Show in sidebar" options and start the add-on.

Once it has started, click on the File editor link in the sidebar menu. The file "/homeassistant/configuration.yaml" should open automatically. If not, click on the folder icon at the top left corner of the panel and select

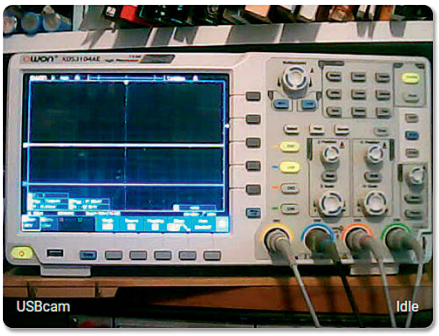

Screen 8: the USB camera's card in the Overview dashboard.

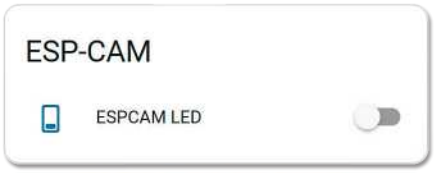

Screen 9: the ESP-CAM card with the LED toggle visible.

"configuration.yaml". Click the red edit icon at the bottom right-hand corner of the screen.

Add the code shown in Block #1 at the end of the file (you need to type the indentation spaces exactly like this, not just the text - indents are two spaces), and save it with the red disk icon at the top right.

Restart Home Assistant using the cog icon at the top right of the screen. Once HA has restarted, the camera card should appear in the Overview dashboard – see Screen 8.

Setting up the ESP-CAM

ESP-CAMs have an ESP32 microcontroller attached to a small camera module. If the USB daughterboard was not supplied with your camera, you will need a USB-to-TTL serial adaptor to complete the process. You can find instructions for doing that from https://pemag.au/link/ac5x

Now connect the ESP-CAM to a USB port on the HA hub. In ESPHome Builder, click on +NEW DEVICE and name your device ESP-CAM or similar. Select ESP32 as the device type. Click INSTALL, then select "Plug into the computer running ESPHome Device Builder".

A USB Serial device should be identified. Select it, and the firmware should automatically compile and upload. It might take some time for the ESP32 platform files and libraries to download before compilation begins. Once the compilation and ESP32 programming are complete, the ESP-CAM should be visible in the ESPHome Builder tab.

Add the code from Block #2 to the device's YAML configuration file in ESPHome Builder and re-install the firmware wirelessly (again, you must include any spaces at the start of the lines exactly like this).

Once the device has rebooted, go to Settings then Devices & services and accept the discovered device, naming it ESPCAM or similar. Re-load the Overview dashboard and a new entry for ESP-CAM, and its associated LED switch should appear within a minute or so – see Screen 9. The ESPCAM LED slider enables the onboard LED.

If the camera image does not appear, it is likely that your camera doesn't follow the original AI-Thinker pinouts. Opening the logs may help you figure out what's going on. Several other pin-out options are available in the documentation at https://pemag.au/link/ac5y

CODE BLOCK #1

```
## USB_camera.yaml
camera:
  - platform: ffmpeg
    name: USBcam
    input: /dev/video0
```

CODE BLOCK #2

```
## ESPCAM.yaml
esp32_camera:
  external_clock:
    pin: GPIO0
    frequency: 20MHz
  i2c_pins:
    sda: GPIO26
    scl: GPIO27
  data_pins: [GPIO5, GPIO18, GPIO19,
  GPIO21, GPIO36, GPIO39, GPIO34, GPIO35]
  vsync_pin: GPIO25
  href_pin: GPIO23
  pixel_clock_pin: GPIO22
  power_down_pin: GPIO32

## Image settings
  name: ESPcam
  id: ESPCAM
switch:
  - platform: gpio
    id: espcam_led
    name: "ESPCAM LED"
    pin: 4
```

CODE BLOCK #3

```
## SamsungTV_IR.yaml
## IR funtions for Pico
external_components:
  - source: github://pr#5974
    components: [remote_transmitter]
    refresh: always
## IR infra red
remote_transmitter:
  pin: 0
  carrier_duty_percent: 50%
## toggle power on/off
button:
  - platform: template
    name: "TV on/off"
    id: TV_toggle
    on_press:
      - remote_transmitter.transmit_pronto:
          data: "paste hex string here"
```

CODE BLOCK #4

```
{{now() - state_attr(''automation.Intrusion'', ''last_triggered'') >
timedelta(minutes = 1) }}
```

CODE BLOCK #5

```
## intruder_alert_enable.yaml
  - platform: template
    name: "Intruder active"
    id: Intruder_active
    optimistic: True
```

These code blocks will be available as a download from siliconchip.com.au/Shop/6/2482

IP camera setup

While integrating an IP camera into HA is straightforward, finding a camera compliant with the ONVIF standard for pan, tilt and zoom (PTZ) can be tricky. Marketplaces like eBay or AliExpress will return hundreds of matches to a search for "ONVIF webcam" or "ONVIF IP camera". Check in the item's specifications for a mention of ONVIF compliance.

I had the most success with V380 cameras that also indicate ONVIF compliance.

Install the camera on your local network using the app that applies to your camera, eg, V380 or V380 Pro. I found the V380 mobile phone app straightforward to use; the cameras

Screens 10 & 11: the IP camera's dashboard card (left), and the dashboard card with pan and tilt controls added.

behave like a WiFi hotspot when first switched on.

This time, we'll launch the installation via a link on the HomeAssistant website. Navigate to the ONVIF page on the HomeAssistant webpage (www.home-assistant.io/integrations/onvif) and click on the blue ADD INTEGRATION TO MY button. Accept the invitation to open another HA settings page as well as the "Do you want to set up ONVIF" prompt.

Allow the device setup to search automatically for your camera. Your camera should appear in the "Select ONVIF device" pop-up options. Select it and click SUBMIT. Name your camera "IP camera" and finish the setup. Refresh the Overview dashboard and the camera should appear, as shown in Screen 10.

The pan and tilt functions can't be enabled in the automatically configured Overview dashboard, as additional code needs to be added to the display card. These functions will be enabled later, when creating a custom dashboard.

A custom dashboard

So far, we've relied on the Overview dashboard, which is maintained automatically by HA. While it displays all the enabled devices, more advanced features are not available. The system supports multiple dashboards, so we will create one to enable the pan and tilt functions of the IP camera and to graph temperature over time.

Go to Settings, then Dashboards and select + ADD DASHBOARD. Select the "New dashboard from scratch" option and make the Title "myDash". "Show in sidebar" should be enabled. Select your new dashboard from the main left menu and click on the edit pencil at the top right corner of the screen.

To create a card for the IP camera in a new section on the dashboard, click on the four squares and a "+" sign icon inside the dashed box to create the section, then click the "+" sign within the section to create a new card. Type "picture" into the search box and select "Picture glance" from the results.

Select your IP camera from the Camera entity drop-down menu. In the "Entities (required)" section, delete the existing entries. Click on SHOW CODE EDITOR immediately below.

In the "entities:" section of the existing configuration, below the existing "- entity" line, paste the contents of the IP_camera.yaml file from the download pack (siliconchip.au/Shop/6/2482) and save the configuration. The "camera_image" line should be after the pasted text. Click DONE to exit editing mode. The pan and tilt controls on the camera card should look like those shown in Screen 11; you can test them now.

If the arrows don't appear on the camera card, it's likely that you have pasted the file into the wrong spot or your camera's stream name is different from the one in the file. In the latter case, make all references to the camera the same as the one in the "- entity" line.

A temperature history chart

This section relies on you having built the Satellite described last month, which incorporates a temperature sensor, relay and LED indicating when the relay is on.

To create another card for the temperature graph in a new section on the dashboard, enter editing mode and click on the "+" sign. Type "history" into the search box and select "History graph". Change "Hours to show" to 1. Under the "Entities" heading, delete any existing entity and add "myHome Temperature" from the drop-down list, then save the result. A graph of the temperature over the last hour should appear, as shown in Screen 12.

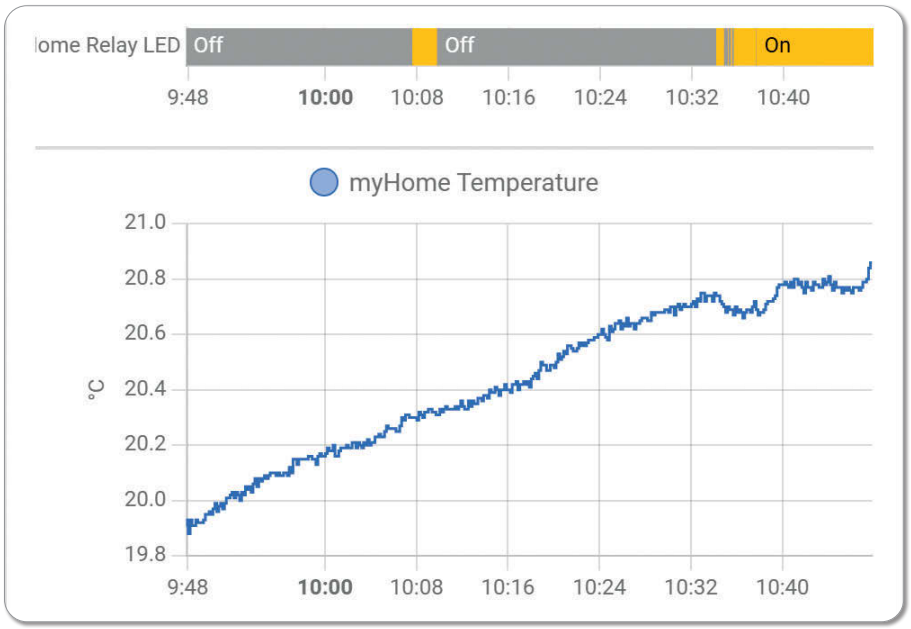

Screen 12: the custom temperature-over-time chart shown at the bottom and the relay LED history shown at the very top.

Home Assistant with a RPi part two

To complete the dashboard, we can add a timeline to show when the thermostat was switched on. Re-open the card by clicking on the pencil that appears when you hover over the control. Add a second entity, “myHome Relay LED” to the card. Click Save and a timeline should appear above the graph. Click DONE to exit editing mode.

Place your finger on the temperature sensor. When the LED lights due to the AC On automation triggering, the end of the grey bar should turn yellow as the temperature graph spikes – see the top of Screen 12.

A better thermostat

The thermostat created in last month’s article lacks a key feature: the ability to easily change the temperature setpoint. HomeAssistant has a generic thermostat integration that can control the relay. First, disable the two existing AC automations in the Automations settings menu by moving their Enable sliders to the left.

Go to Devices settings and add an integration called “Generic thermostat (helper)”. Name it “AC Thermostat”; select “Cooling mode”, use the “myHome Temperature” sensor and “myHome Relay LED” as the Actuator switch, then set the Cold tolerance to 1. The minimum and maximum target temperatures can be set to 10 and 30, respectively. Scroll down and click NEXT.

Add temperature values into the Comfort and Eco presets, then click SUBMIT and FINISH.

The card will appear in the Overview dashboard – see Screen 13. The setpoint is displayed in the centre of the card and is controlled by moving the larger open circle on the gauge. Enable the thermostat by clicking the snowflake icon. The relay LED should switch on when the setpoint is 0.5°C lower than the room temperature, and off when the setpoint is 0.5°C higher.

The activity bar above the temperature graph that we created earlier should show the Relay LED’s activity as the thermostat is exercised.

To use one of the temperature presets, click on the three dots at the top-right of the card. A drop-down menu at the bottom right of the pop-up allows selection of any available presets (Screen 14). Selecting None will revert to the last manual setting.

A home air conditioner (especially a split system) is likely to have an infrared remote control. The TV remote control below could readily be adapted for this purpose.

IR remote control

Adding remote functions for devices with infrared remotes involves adding a button to your dashboard and connecting it to an appropriate code for the IR LED to transmit (like the one on our Satellite).

Finding the correct codes for your device may take some hunting around. The most comprehensive source I have found is IRDB (https://pemag.au/link/ac5z). The free sign-up allows five codes per day to be downloaded.

The code begins with a patch to enable the IR function on a Pico. By the time this is in print, the “external_components:” section may no longer be required. If the code compiles properly, then leave it in; otherwise, try without it.

A Button input is then defined, which will appear on the dashboard. Finally, get the PRONTO IR code from IRDB and paste it into the “data:” line. A separate section of code starting with “– platform: template” will need to be created for each additional function button (see Block #3).

The SamsungTV_IR.yaml code in the download pack includes the code to toggle a Samsung TV’s power. Once the Satellite has been updated, a “TV on/off” button labelled PRESS should appear in the myHome card on the Overview dashboard – see Screen 15.

If you have a Samsung TV, point the IR LED toward the TV and press the button. The TV should switch on or off with each button press. The appropriate codes for many other TV brands should be available from the IRDB database.

Access from a phone or tablet

To provide access to a smartphone or tablet dashboard while connected to your local WiFi network, install the HomeAssistant app for iOS or Android.

Start the app & follow the prompts. If you don’t enable location permissions, the device tracking function and the sensors on the mobile device won’t register properly with HA. The Notifications permission will be used in a later example, so enable it now.

Enter the username and password you created earlier in HomeAssistant into the app. All HA functions should

Screens 13 & 14: the thermostat card (above) and accessing the thermostat presets (below).

Screen 15: myDash on a mobile phone, with the TV on/off button visible near the bottom.

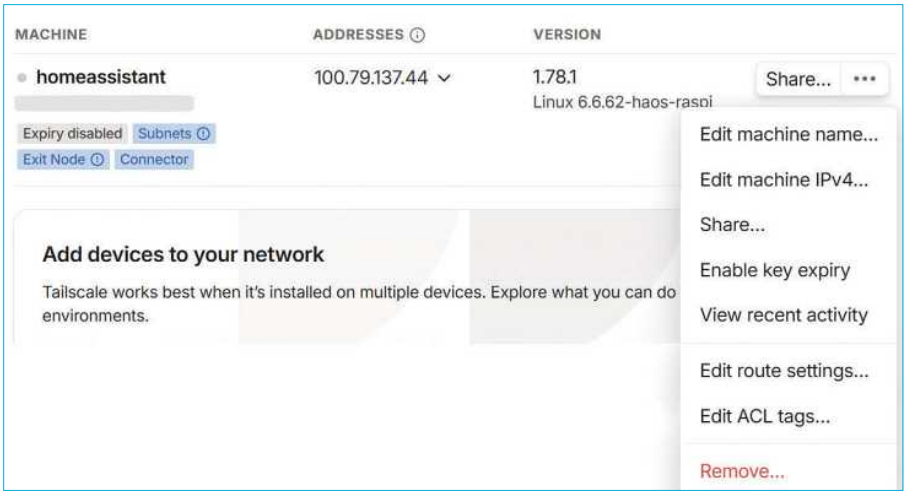

Screen 16: the TailScale website console page. The HomeAssistant entry needs to be made an Exit Node, and the detected subnet approved in the "Edit route settings" dialog.

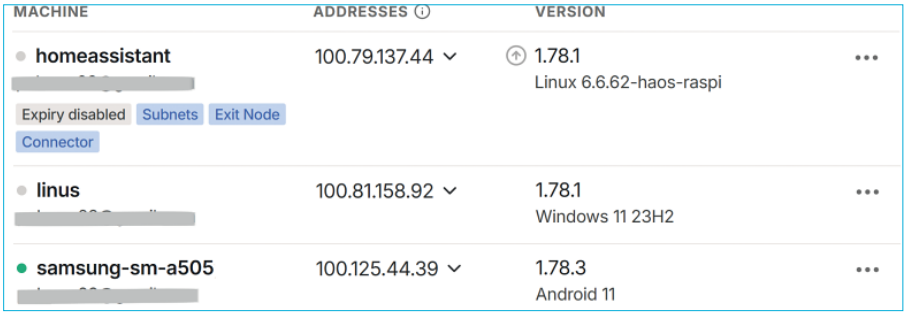

Screen 17: the TailScale control panel after a PC and phone have been added.

be available. As mentioned before, creating and using a profile on HA without the ability to administer the system is a good idea for everyday and particularly mobile use.

Virtual private network

Remote access to HA from outside your local network requires several elements: a URL that is recognisable externally, a way to convert that to an IP address and secure access to the HA hub through your internet router.

There are several different approaches to providing remote access to HomeAssistant, which are discussed at https://pemag.au/link/ac60

WunderTech provides a complete examination of the pros & cons of each method at https://pemag.au/link/ac61

Virtual Private Network (VPN) connections are relatively easy to set up and offer a high level of security. TailScale has a free offering that is a recommended HA option for remote access.

TailScale needs to be installed on every device that has remote access. Each connects to TailScale's cloud service to discover the other members. In most cases, no router settings need to be changed when using TailScale.

The VPN requires setting up a TailScale account plus configuration of HA and your mobile device. If you find the instructions below difficult to follow, there are several good YouTube tutorials on connecting HA to TailScale. I found the one by Joyce Lin to be helpful (https://youtu.be/EJ3cjoJAaQA).

Go to TailScale's website (https://tailscale.com) and set up a free account. It requires you to use an existing identity provider rather than setting up a new username and password. If your preferred identity source is not in the list, the OpenID Connect (OIDC) option can provide linkage to a wide range of additional providers.

To add the integration in HA, go to the Settings then Add-ons menu and click the + ADD-ON STORE button. Search for "TailScale" and follow the installation prompts. Enable all four options in the configuration panel, then start the add-on and then click the OPEN WEB UI button.

Log HomeAssistant into TailScale with the same credentials you used to create the account, re-authenticating if required. Next, connect your HomeAssistant to your TailNet.

The final step of the installation process on HomeAssistant opens a browser window that logs you into TailScale's control panel. Under the "homeassistant" entry, there will be several blue indicator boxes – see Screen 16. If they have exclamation marks in circles, click on the ellipsis (...) at the end of the homeassistant line and select "Edit route settings". Select both check boxes and then Save.

The "Exit node" option allows HA to act as a server on your personal TailNet.

You can now log into HomeAssistant from any device that is connected to your TailNet. The Subnet router option authorises the HA hub to route traffic from satellite devices through the TailNet.

Install the TailScale app on your mobile device from the app store and log in to it. The HA app should now be able to access your hub when outside your wireless network's range. Similarly, TailScale can be added on a PC – see the "linus" entry on Screen 17.

Switching off WiFi on your phone, if it is enabled, will allow this feature to be tested. The HomeAssistant app on your phone will ask for a new IP address. http://homeassistant:8123 should do the trick, as TailScale provides address translation and routing services. After a few seconds, your connection to HA should be restored.

An alternative, configured with just a few mouse clicks, is Home Assistant's Cloud service, which has a 31-day free trial followed by subscriptions at around $11 per month or $110 per year, including GST.

Your default HomeAssistant login has full administrator privileges on the remote device. For better security, creating a second, less-privileged account. Go to the Settings then People menu on the main HA screen then click + ADD PERSON and it will lead you through the process. Do not enable the Administrator option for this new account.

Thermostat notification

Notifications are implemented by creating an action in an automation.

To create a new automation, add a Trigger, selecting Entity and then State as the type. In the When panel, select "AC Thermostat" as the entity, "Current action" as the Attribute. Set "From" to Idle and "To" to Cooling, then click on + ADD ACTION at the bottom of

the screen, select “Notifications” and “send a persistent notification”

Fill in an appropriate message and click the SAVE button – see Screen 18.

Each time the virtual AC switches on, a timestamped notification message will appear in the Notifications panel on all connected devices.

To disable the automation, go to the Automations menu & move the slider attached to the automation to the left.

A PIR intruder alert

Now let’s create an automation that produces a notification each time the PIR sensor on the Satellite is triggered. The trigger will be the myHome.PIR entity. Set “From” to Clear and “To” to Detected, then click on “+ ADD ACTION” at the bottom of the screen and select Notifications.

If you have set up remote access on a phone, select “send a notification via mobile_app_xxxx”. Put “Intruder alert!” in the message and in the title and click SAVE; name your automation “Intrusion”. If you have not enabled remote access at this point, use the “Persistent notification” option instead.

Every time something moves in front of the PIR sensor, a notification message will be sent. If a human intrusion occurs or the cat decides to play in front of the sensor at midnight, you’ll be flooded by messages. A time delay on the automation being re-triggered will limit the number of messages.

To delay re-triggering, a condition is added to the automation that detects the last time it was triggered and calculates whether it was longer than one minute ago. The automation is not re-triggered if the time was less than this.

Edit the Intrusion automation, using the “And if” clause, and select Template from the “Other conditions” option. Enter the code in Block #4 (all on one line), which is also in the intruder_alert_enable.yaml file, into the “Value template *” box.

Save the automation and test it. You should be able to trigger the PIR sensor multiple times, but only receive a notification on your mobile device if the PIR is triggered more than a minute after the last notification.

To create a switch to turn off these notifications entirely, place the code from Block #5 directly under the Relay LED switch code and update the Satellite.

Edit the Intrusion automation, adding another condition. Select Entity and “myHome.Intruder active”. Set the State field to On, then save the automation. The new switch should appear on the myHome card in the Overview dashboard – see Screen 19. If the switch is off, no notifications should be sent.

Two useful tools in debugging automations are the Logbook in the main menu and TRACES when editing the automation. In TRACES, clicking on an icon in the sequence will display what happened at that point in the automation’s flow.

Backups

Backups are important in maintaining the integrity of your setup over time. HA’s update function triggers a local backup of specific components each time the configuration is changed, or the system is updated. Backup information is stored in several folders on the SD card.

To enable regular backups, go to Settings then System then Backups. In the “Backup settings” panel, enable “Use automatic backups” and adjust the Schedule and Retention parameters as desired.

In the “Backup data” panel, disable History and Media, as these can generate large video files that may fill the SD card and crash the system.

Scroll down to the “Encryption key” panel and download the emergency kit. The encryption key and instructions in the kit will be required to restore from the backup.

If you want to save backups externally, HA can access external and cloud storage. I found that the most straightforward path was to make HA’s files accessible via an SMB (CIFS) file share (see https://pemag.au/link/ac62).

To do this, install the “Samba share” add-on. Using the Configuration tab at the top of the installation page and add a password for remote access. The Username and Workgroup fields can be left as they are or changed to meet your needs. Start the add-on with “Start on boot” and Watchdog both enabled.

After a minute or so, you should be able to see a HOMEASSISTANT host in your computer’s file browser under Network (if using Windows) or the appropriate heading on Mac or Linux (you may need to install the samba package to access it in Linux).

If it doesn’t appear, try refreshing or re-opening your file browser window. Map network drives for the “backup” and “config” folders, using the credentials you saved in the Samba configuration tab. Use your PC’s backup tool to regularly save the files.

Conclusion

In this series, we have touched on some key elements of home automation using the HomeAssistant and ESPHome platforms. We hope it has whetted your appetite to explore more of the many features these platforms offer. PE

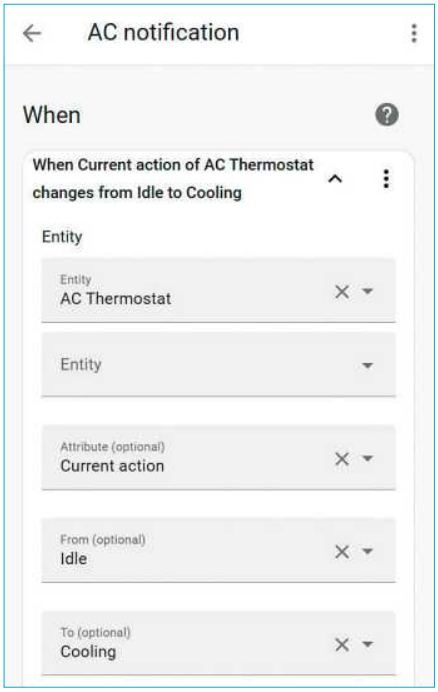

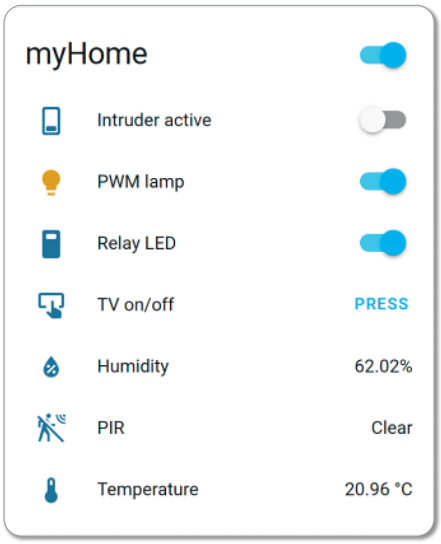

Screens 18 & 19: setting up the “When” parameters for the AC Thermostat notification (left). The myHome Overview dashboard card with all features enabled (below).

AUDIO OUT

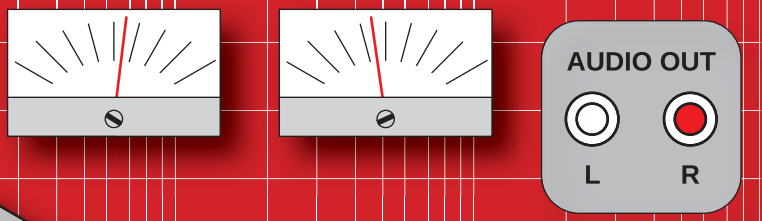

By Jake Rothman

The RetroTech UK event and radio rallies, part 2/2

Last month, I explained that I had recently attended RetroTech UK, an event run by the British Vintage Wireless Society (BVWS – bvws.org.uk). It's a wonderful sale of old electronic stuff that takes place every May.

I then described many of the fascinating analogue and vintage objects I came across at that show, ending with some interesting valve-based TV cameras before I ran out of space.

I shall now pick up where I left off, looking at some more old TV cameras, followed by some other items I came across.

More TV cameras with tubes

The website tvcameramuseum.org also has some good data on lots of other cameras. The monochrome camera shown in Photos 22 & 23 used 91 valves and consumed 1.1kW. Consequently, the power supply was a sight to behold, glowing away in Photo 24. Brian also pulled out the hard-wired chassis, shown in Photo 25.

Satellite truck

Photo 26 shows another exhibit from the Broadcast Engineering Museum, a satellite outside-broadcast van. It was only 20 years old and is now history. The interior was even more fascinating, as shown in Photo 27.

Distribution amps

The museum had an area at the front of the show where they were selling off their surplus bits 'n' bobs, all totally unique. I bought a quadrant-fader camera control unit, some diecast boxes with weird BBC one-off circuits, and another Mosfet power amp.

Photo 22: having a good confab with Brian Summers, who restored and exhibited the classic Marconi TV camera.

Photo 23: inside the camera; all metalwork and valves. I've got the relays in my stock. I'll send them some.

Photo 24: the camera's power supply valves. The two with the pink glow are gas-discharge voltage regulators.

Photo 25: pulling out the hard-wired chassis in the control unit. I was amazed to see that most of the components were the originals. Only a few resistors had gone open circuit and had been replaced

Photo 26: an old (20 years!) outside broadcast satellite-link van.

Photo 27: the van's interior. It had a TSL sound monitor unit in the lower left-hand side with the two PPM meters. We use one in our studio; a fantastic bit of design.

Photo 28: an Avitel video distribution amplifier rack. It's worth more than that just for the metal frame.

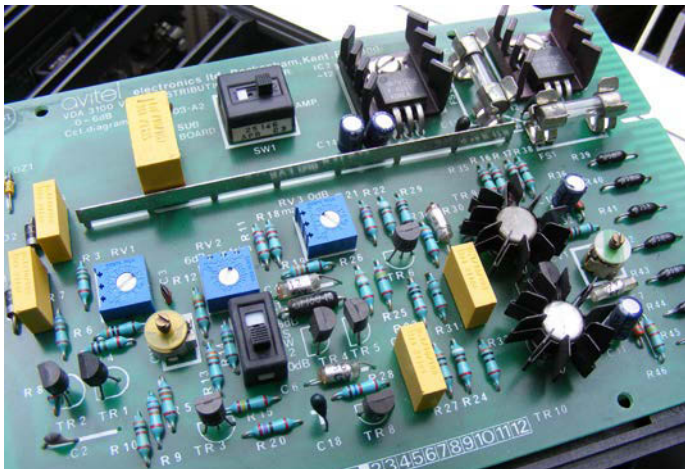

Photo 29: the PCB shows the excellent idea of a soldered-in thick metal ground bus.

More interesting was a rack full of Avitel video distribution amps (see Photo 28). I originally bought this for the rack unit to use for Eurorack synths, but when I got home and pulled the cards out, as shown in Photo 29, I was intrigued.

Note how the components are well-labelled; no printing under the components on this board!

As well as audio power amps, some of the last bastions of discrete circuitry were video distribution amplifiers. The circuit is basically a high-speed op amp with a high-current output stage designed to drive six 75Ω loads.

This would make them excellent headphone amps, although this depends on the noise and distortion levels, which may be compromised since video amps are optimised for bandwidth, which needs to be at least 5MHz. It would also be good for a signal generator output stage. We can't have an *Audio Out* without a circuit and its description, so it's shown in Figs.1 & 2.

Fig.1: the Avitel video distribution amplifier JFET input stage. I like discrete circuitry; it's like low-level coding, you have total control over every step.

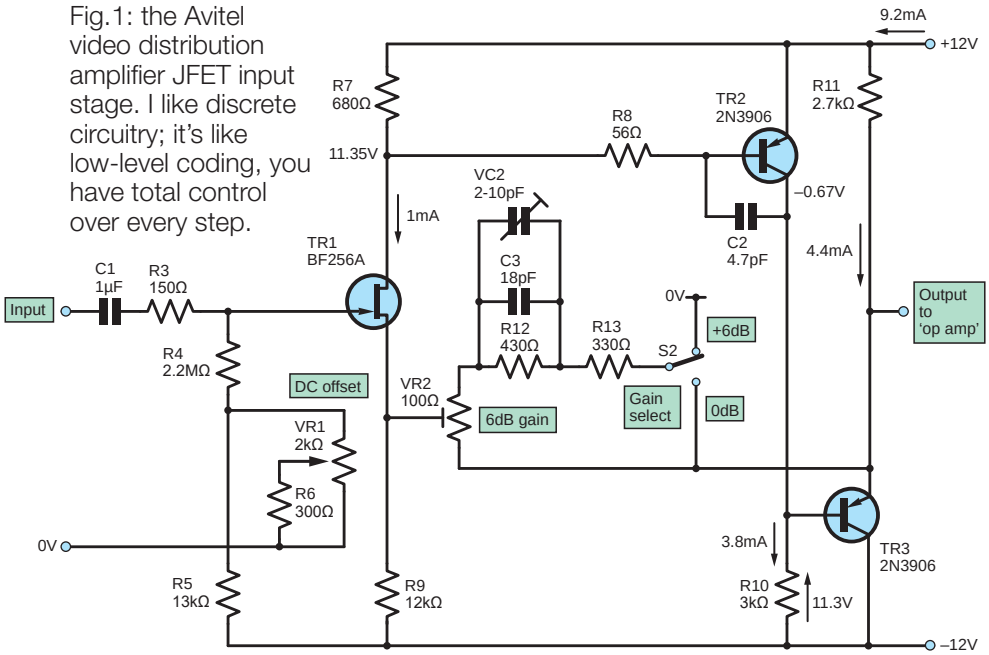

Starting at the beginning, we have a JFET input buffer for a high 2.2MΩ input resistance set by R4. The JFET stage, TR1, has a small amount of gain, in the order of 5×. This is then multiplied by the higher gain of the common-emitter stage based on TR2, which would be around 50×. Output buffering is provided by an emitter-follower, TR3.

The open-loop gain is brought down to unity via negative feedback through R12, R13 and R14. There is a +6dB boost function selectable via S2, trimmable by VR2. A variable DC offset-null function for the whole amplifier is provided by VR1.

We then move on to the 'op amp' with a long-tailed pair built around TR4 & TR5. This has a response-peaking circuit built around inductor L1 and capacitor C4, damped by

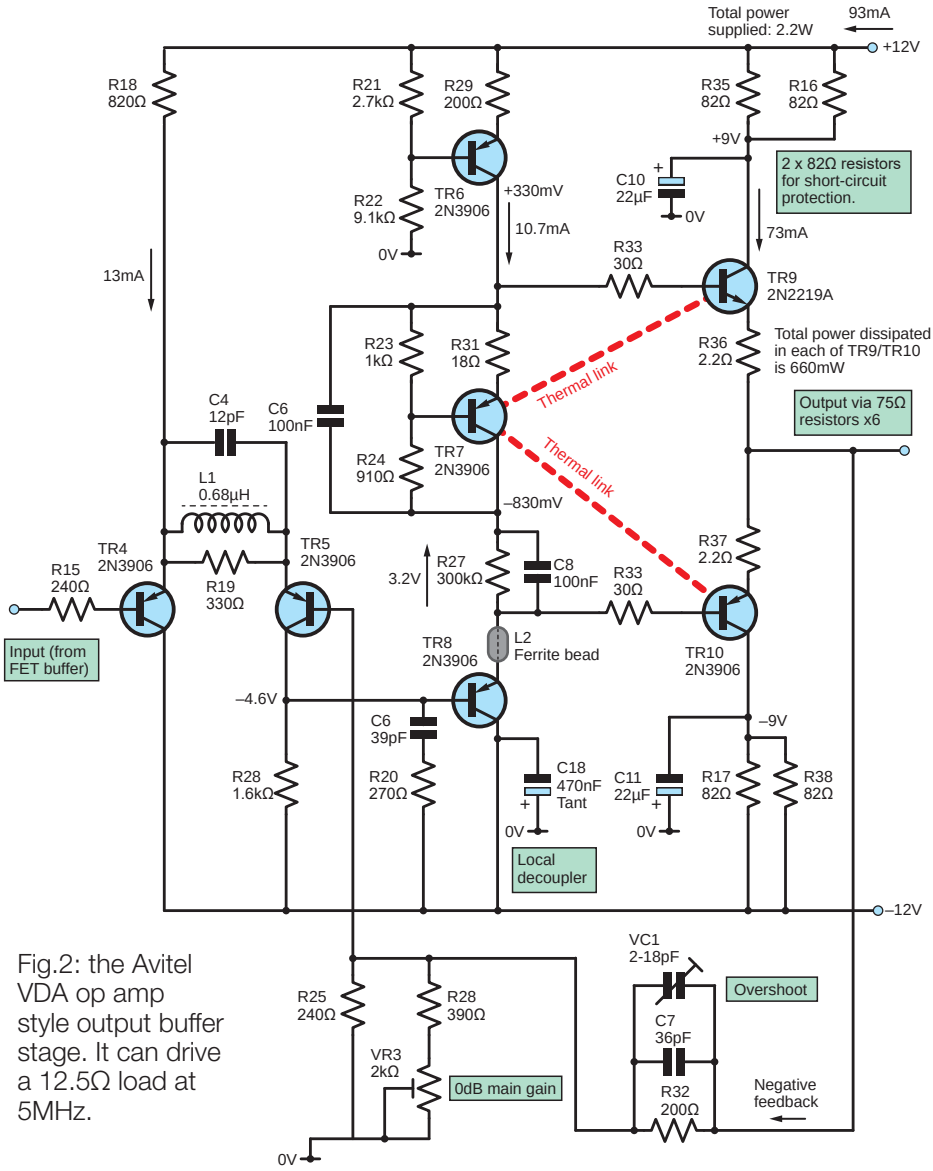

Fig.2: the Avitel VDA op amp style output buffer stage. It can drive a 12.5Ω load at 5MHz.

resistor R19. The long-tailed pair (LTP) then feeds an emitter-follower built around TR8. A speed-up network is formed by C8 and the ferrite bead shown dotted.

TR8 is loaded by a constant-current source built around TR6. Note that the currents through the LTP and the follower are quite high at 13mA and 10mA, respectively, to give maximum bandwidth.

Unusually, all the small-signal transistors are of the same polarity, all PNP. This reduces mistakes in hand assembly and keeps the inventory down.

The output stage is a Class-A push-pull complementary follower with a pre-set standing current of 70mA. This is stabilised by transistor TR7, which is wedged between the heatsinks of the output transistors, TR9 and TR10. Note these are relatively fast 300MHz TO-5 package devices, rather than the slower TIP31/2 3MHz types that would typically be used for audio work.

There is little open-loop gain in this circuit, only that provided by the LTP at around 40×. That means the distortion will be fairly high.

There were some more modern VDA cards in the rack with DC restoration, but they used ceramic thick-film modules, making them difficult to experiment with. The construction of this unit is an excellent example of sound, reliable 1980s design. Each module had individual fused on-board regulators.

Avitel is still in business, but I didn't have any luck finding a circuit. I later contacted Dave Hill (BECG website manager) at the museum, and he told me they had all sorts of data in their archive, and he found it for me.

Analogue meters

These shows are the place to find analogue meter movements. I like the BBC-developed peak programme meters with their logarithmic curve and 1-to-7 scale markings. This is an analogue instrument that has found a new lease of life in digital recording.

I worked at a small company called Soundex that made driver boards for Sifam PPM movements. I was pleased to find an old BBC rack panel for £45 with left/right and mid/side dual-concentric pointer meters using Alice PPM704 cards, shown in Photo 30. Yes, there will be an article with some inside information on these unusual piece-wise-linear logarithmic circuits.

Computers

Old computers are not really my

Photo 30: an Alice dual stereo PPM meter. £45.00 at RetroTech 2026, £810 at Canford Audio in 1996.

Photo 31: digital buffs are also catered for. I love orange CRT monitors before bed. No insomnia-inducing blue light.

Photo 32: a pair of lovely coffee table books, one on KEF loudspeakers and one on Quad.

Photo 33: some 1980s high-tech audio test gear from Wayne Kerr and Brüel & Kjær.

thing, but I couldn't help feeling twinges of nostalgia looking at old BBC Micros from the early eighties (Photo 31), developed when the BBC was biased in an educational direction. There was even a Sinclair ZX-81, on which I wrote my first useful program for calculating power supply outputs.

I remember upgrading it to 16kB of memory with a stack of chips hanging off the back. They usually fell off at some vital point.

Also present was an old Sinclair calculator kit, like the one I forced my dad to buy for school. Guiltily, I never used it for A-level physics because it had idiosyncratic Reverse Polish Notation. I defected to a Casio because it had a clear green fluorescent display. I still use it today because I know where the buttons are without looking and it uses AA cells.

Books

It's said everything is online; no, it isn't. There is a lot of printed material that's never been scanned and never will be because the owners are non-digital and it will go in the skip when they die. This has already happened to a few top audio engineers.

A large print library of offline material is an essential resource for any designer. I always come back with books and service manuals from these shows (Photo 32). They are not quite the same on a Kindle; expensive, though, at £45 each, but I don't skimp on my library.

I've bought quite a few beautifully drawn Heathkit manuals for £3.00 recently, since these have been locked up behind a paywall by some copyright-hoarding rentier. We need wealth creation, not extraction!

I was told once we'd make more money knocking down a North West London electronics factory I worked at in 1988 and making it into a car park. The consultants were right, of course.

Photo 34: a stack of Uher tape machines, a small reel-to-reel machine used by local radio.

Test equipment

There's no need for a budding audio engineer to buy new bench power supplies, oscilloscopes or signal generators; there's always plenty available at these shows. Most of the faults are mechanical: resistors gone high in value, and dried-out electrolytic capacitors.

Some test gear can be difficult to fix and calibrate, though, such as those shown in Photo 33 that often have burnt-out CRTs. I have a friend who is storing ex-NATO Gould OS3000 oscilloscopes, waiting to wind a viable quantity of new mains transformers for them.

I've been working on restoring a Wayne Kerr AMS1 audio distortion analyser for the last three years. It's worth it, though, because it doesn't need software or a PC.

Tape machines

I need a new cassette deck; I've been waiting for a suitable one to turn up with no luck. Most of the ones I've seen needed new heads, unobtainable belts, or new rewind cogs. There were plenty of reel-to-reel machines, such as the ex-local-radio Uher machines in Photo 34.

The Revox B77, with its three separate motors giving good reliability,

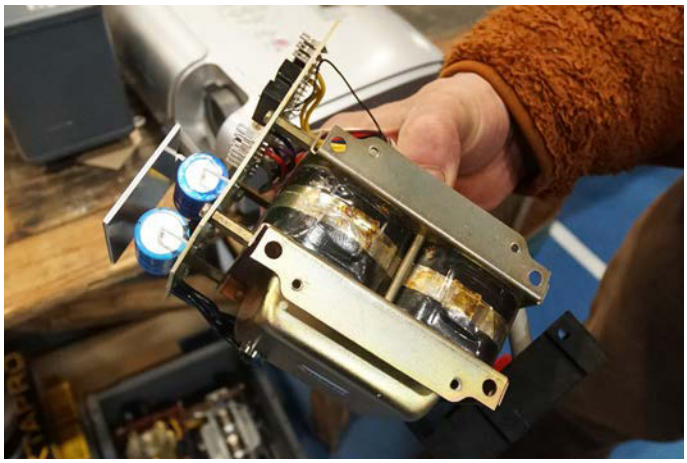

Photo 35: an unusual C-core power transformer for a Revox B77.

Photo 36: the only music gear stand at the show. No, there wasn't a Roland TR808 drum machine for £20.

is especially well catered for. There were a couple of stands full of spare parts (see Photo 35). You can get some weird spares at these shows, just don't go expecting to find what you want; you'll find something you forgot you needed instead.

Music gear

There's usually very little old electronic music gear, but due to recent digital bloat build-up on platforms such as eBay and Reverb, it has been increasing as people cut out the middleman and the courier. Photo 36 shows a stand this year.

Don't buy a Focusrite interface like the one shown in Photo 36; they have proprietary software that screws your computer up. Hopefully, there will be more music stuff next year.

Döstädning

No, I'm not swearing, nor is this the name of the latest Ghost single. Still, one is expected to end on some miserable Scandinavian note today; in this case, the Swedish death cleanse.

It's worth remembering that a square-foot pile of analogue electronic junk may only be worth around £20. Contrast this with the value of real-estate square-footage and the answer is obvious: clear it out. I'm already feeling the pressure.

Two Transit van loads constitute a typical domestic electronic workshop, for which a house clearance company would charge £400, even if the stuff in it was originally valued over £10,000 on eBay (dream on).

Much of the wonderful cheap stuff at these shows comes from deceased persons, described aptly at radio rallies as "silent key" sales, referring to their Morse code key. There should be an audio equivalent: "mute speaker" or "cold iron", perhaps?

I'll get my first stand at RetroTech next year, so I'll see you there. My family suspects I'll come back with more than I go with.

"Ooh" (to use a Welsh expression I've apparently picked up), "that RDG Radiogram looks interesting!" It was as well (see Photo 37). I was very pleased to find it was stereo and had two EL84 single-ended valve amplifiers. The only problem was that it had only one speaker. There was supposed to be another matching four-legged speaker that you could add later on.

It only cost £20 and had a 10-inch (25cm) Celestion speaker in it as well. I'm going to restore its two valve amps for an *Audio Out* column. I won't bother with the record deck though, it would destroy my records with its huge, heavy arm.

Making it middle and side, so stereo is emitted from the one cabinet, seems a good idea. Radiograms seem to be making some sort of comeback. I wonder how the man in Photo 38 managed to co-opt his wife into carrying one!

Many thanks to my son, Harvey Rothman, for taking the photos and to the traders who were happy for us to take them. PE

Photo 37: this big valve RGD gramophone was left at the end of the show; my son Harvey got it for £20.

Photo 38: my wife wouldn't help me hump electronic junk into the car. What's his secret?

AOShop

The home for specialist audio, analogue and historic components – provided by Jake Rothman, *PE's Audio Out* columnist. The AOShop is your best bet for classic analogue 'NOS' (new old stock) components, including all parts for *Audio Out* projects and designs.

Ge semiconductors

Small-signal PNP transistors
NKT214F, OC57, OC59, 2N1377, 2N525 **£1.00**

Low-noise PNP transistors
GET106 **£2.50**

Small power PNP transistors
AC153, AC153K, AC188K **£1.50**

Small power NPN transistors
AC176, AC141K **£2.00**

High-voltage PNP transistors
OC77, CV7001 **£1.50**

TO3 PNP power transistors
AD149, AD161, AD162 **£2.00**
AD140 **£3.50**
AD149, AD161/2 matched pair **£5.00**

RF PNP transistors
OC41, OC42, 2SA12, 2SA53, AF124, AF178, GET872A **£1.50**

NPN transistors
OC139, OC140, ASY73 **£2.00**

Diodes
CV7049 (OA10), CG92 (OA91) **£0.50**

Si semiconductors

Diodes
ZC5800 RF Schottky **£0.20**

Low-noise silicon transistors
BFW16A, 2SC3071, 2SC3068, 2SA1016K, 2SC2362K, 2SA970BL **£1.50**
2SC2204, 2SD655, BC550C **£0.50**
ZTX651 **£0.30**
ZTX751 **£0.50**

RF transistors (suitable for Theremin)
BF180 **£0.40**
BF199 **£0.50**

Audio power MOSFETs
Exicon 10N20, 10P20 **£6.50**
Hitachi 2SJ99, 2SK343 **£3.50**
Hitachi 2SJ56, 2SK176 **£8.50**

JFETs
BFW11, BFW10, TIS73L, J177, J113, U1994, U1898, 2SJ176, J201 **£1.00**
J175, J176, J112, J111, 2N3820, 2N5457, BF244, 2N5460, J230 **£0.60**

Small power output/driver transistors
2SB649A, D669A, 2SA1208, 2SC2910, MJE253G, MJE243G, 2SA1725, 2SC4511 **£1.20**
BD139,BD140, BD135, BD136, BD435, BD436, MPSA63, BCV46 **£0.50**

MOSFETs
ZVP2106A **£0.30**
Dual-gate 3SK45, BFS28 **£1.50**

Metal-cased transistors
BCY71 **£0.30**

Amplifiers
LM318 high-speed op amp **£0.35**
µA709 metal-cased op amp **£2.00**
5534H metal-cased op amp **£3.00**
LM384, TDA2030A, TDA2050V **£1.50**
TL082 (ceramic pack) **£1.00**
TAA435 (Mullard power amp driver for AD161/AD162) **£1.20**

ICs

That/dbx 2180 VCA/VCF **£6.00**
CA3080 VCA/VCF (vocoder) **£2.80**
CA3280 VCA/VCF **£4.00**
LM13600/ LM13700 VCA/VCF **£2.00**
PT2399 echo/delay **£1.50**
Original TI 44V NE5532 **75p**
Nisshimbo 44V NE5532DD **£1.00**
SMT version **50p**
New 36V TI NE5532 **50p**
Original TI NE5534A **£1.40**

Dual transistors

2N2639, 2N2223, 2N2910 (NPN) **£4.00**
E401 (JFET Moog) **£4.00**
2SK2145-Y dual JFET **£0.80**
2N5564 JFET **£8.00**
HN3C51F, HN3A51F **£1.00**
DMMT3904/6, HN1A01F, HN1C01F **£0.50**

Loudspeakers

PE Mini-Monitor
Volt PE165 6.5-inch woofer (each) **£85**
Kit *pair* of PE165/MDT29 plus Volt crossover parts and PCBs **£299**
Monacor DT-28N tweeter (each) **£35**
Vifa 19mm BC20SC15-04 tweeter (each) **£15**

Volt crossover inductors
1.2mH, 1.5mH, 2mH, 2.7mH, 0.5mH (tapped at 0.3mH) (each) **£5.85**

Fully tested reclaimed speakers
Vifa BC14 5-inch woofer (each) **£10**
Vifa TC26 1-inch tweeter (each) **£10**

Low-price speakers
Philips 4-inch 4070 **£2.00**
EMI 10x6-inch, 30Ω Alnico **£7**
64mm 64Ω neodymium **£1.20**

Capacitors

Note '10/63' denotes '10µF 63V'.

Polyester
3.3/100, 4.7/250, 4.7/63 **£1.00**
5.6/63, 8.2/63, 10/63 **£2.00**
Mullard 'Mustard' C296 0.22/400 **£2.00**

Polycarbonate
Axial 2.2/63 1%, 4.7/160, 6.8/63 **£1.00**

Radial
6.8/160V, 10/63 **£2.00**
Reclaimed 22/63 **£2.00**

Polystyrene
Philips 1% 4.7nF/160, 6.2nF/500, 12nF/63, 22nF/63, 110nF/63, 24nF, 2nF **£1.00**
RIFA 2% 100nF for active xov. **£1.00**
Suflex 90.9nF 0.5% **£2.00**
Suflex 2.5% 10nF/63 (rad. or ax.) **£0.50**

Radiation resistant
Siemens cellulose acetate MKL 2.2/25 **£0.80**

Electrolytic – Mullard blue 017 series
10/25, 22/25, 100/10 **£0.50**
150/40, 470/40, 1000/25 **£1.00**

Tantalum – axial metal cased
22/50, 47/35, 68/25, 100/20, 120/10, 150/16, 220/10, 330/6 **£2.00**
22/35, 33/35, 47/20, 68/15, 100/10, 150/6 **£1.25**
4.7/50, 6.8/35, 10/25, 10/35, 22/15 **£1.00**

Axial moulded-case tantalum
Kemet axial 6.8/10 **£0.30**
Kemet radial 33/10 **£0.40**
STC radial 100/20 **£1.50**

Tantalum bead
22/50 **£1.00**
470/3 **£2.00**
680/6.3 **£3.00**

Wet tantalum
220/25 axial **£3.20**
Castanet button 140/30, 470/3 **£3.20**
Hughes 540/10 **£3.20**

Bipolar
Hermetic bipolar tantalum 16/35 **£3.20**
Elcap axial 10/50 **£0.50**
Generic radial 100/16, 470/35, 100/ 63, 22/35, 4.7/35, 220/16 **£0.50**

Philips solid-aluminium (axial) 121/123
47/16 **£1.00**
330/6.3 **£2.00**
100/35 **£4.00**

Philips Pearl 122 series (radial)
10/16 **£0.30**

Silvered mica
1nF/125 axial 1% **£0.30**

Trimmer capacitors
Vishay plastic-film 4-40pF **£0.80**
Vishay plastic-film 5.5-45pF **£1.00**
Vishay plastic-film 5-80pF **£1.20**
Vishay plastic-film 10-250pF **£2.00**
Mica 1-12pF, 2-40pF **£1.00**

Audio transformers and inductors

Eagle transformers
LT44, LT722 driver, LT700, LT723
500Ω output **£2.50**
LT30 500mW output **£3.50**
5:1 interstage (fuzz box) **£1.50**

Repanco
T/T3 splitter transformer **£4.00**
CH2 5mH RFC **£2.00**

Balanced output transformer
Vigortronix 600Ω VTX-101-007 **£10**
Vigortronix 600Ω VTX-101-3001 **£10**
Vigortronix 600Ω VTX-101-3002 **£15**
Gardners 150Ω **£10**
Reclaimed BBC LL74/MPC nickel core 600Ω **£12**
Low-level 1:1 600Ω (modem) **£2**
Reclaimed mic input transformer **£15**

Inductors
82µH, 4.7mH, 100µH, 270µH, 10µH, 14µH (low Z) **£0.50**
4-inch ferrite rod with MW and LW windings **£3.00**

Special resistors

Bourns wire-wound trimmer
5kΩ 3059 JM panel-mount **£2.00**
Thermistor RA53, R13 **£4.00**
A13 **£2.00**
Thermistor CZ1, CZ6 **£1.50**
Holco H2 2.2MΩ 1W, 1% **£1.00**
Welwyn 1GΩ 2W **£1.00**
5k Bourns 3321H cermet trimmer **£0.50**

Potentiometers

Bourns 81
25kΩ lin cermet **£2.00**
25kΩ lin, 5kΩ lin cond. plastic **£3.00**

Bourns 91
10k dual-gang lin or log **£10.00**

Plessey moulded-track
5kΩ log with switch **£3.00**
50kΩ A/log **£3.00**
Mil 250kΩ lin dual **£5.00**

Bourns
8-gang 50k A/log **£7.50**
ALPs 4-gang 100k A/log **£3.00**

Alpha
16mm 4.7kΩ A/log **£0.80**
220kΩ A/log **£0.80**
10kΩ lin centre-detent dual-gang **£1.50**

Allen Bradley J series/Honeywell
10kΩ lin 1 million cycles **£5.00**
Blore Edwards AB 45 dual 5kΩ A/log with switch **£3.50**
Alps RK9 dual-gang 5kΩ RD law **£4.00**
Alps 50k log tapped motor. stereo **£5.00**
BI P260 500kΩ log conductive plastic 1 million cycles **£2.00**
Colvern wire-wound 100kΩ or 50kΩ dual-gang 3W **£5.00**
Mil Spec hermetic 10Ω **£8.00**

Miscellaneous

PCBs
Phase-linear crossover PCB **£3.00**
8-gang All-pass Filter PCB **£3.00**
8-gang Pot PCB **£1.50**
Polarity Inversion PCB (Alps) **£1.50**
Polarity Inversion PCB (alt.) **£1.50**
All-pass Filter PCBs (three) **£4.00**
.. Rotary Switch (AO0626ROT) **£1.50**
.. Pot Board (AO0626POT) **£1.50**
Silent Stomp Switch (pair) **£4.00**
Penta-fuzz Bender (*PE*, Feb26) **£3.00**
Pre-driver (*PE*, Jan26) **£3.00**
Pre-driver jack bd (*PE*, Jan26) **£1.50**
Fuzz Box (*PE*, Nov25) **£3.00**
Fuzz Box jack bd (*PE*, Dec25) **£1.50**
Dual JFET Buffer (*PE*, Oct25) **£2.00**
Mosfet Buffer (*PE*, Oct 2025) **£2.00**
Dual Mosfet Buffer (*PE*, Oct25) **£2.00**
Cascode Amp/PS (Sep 2025) **£2.00**
Pocket Theremin (*EPE*, 1996) **£2.00**
Elysian Theremin (*EPE*, 1996) **£6.00**
Synth VCF, VCO (*EPE*, 2017) **£3.00**
48V PSU (*EPE*, 2019) **£3.00**
SMT dual transistor adapter **£0.40**

Switches Blore Edwards/Plessey/etc
ALPs 4-pole push switch **£1.50**
MH1 4-pole 3-way **£3.00**
Elma 01 2-pole 6-way **£6.00**
3PDT guitar pedal stomp switch **£3.50**

Valves
EC88, PCL84 **£3.00**

Contact
Jake Rothman
Manchester House
Market Street, Craven Arms
Shropshire SY7 9NN
(visit by appointment)
+44 (0)1597 829102
jrothman1962@gmail.com

Payment
PayPal, bank transfer, cheques (payable to "J Rothman", UK pounds only). No VAT payable

Postage
Small Jiffy bag £2.99
Small package £4.99

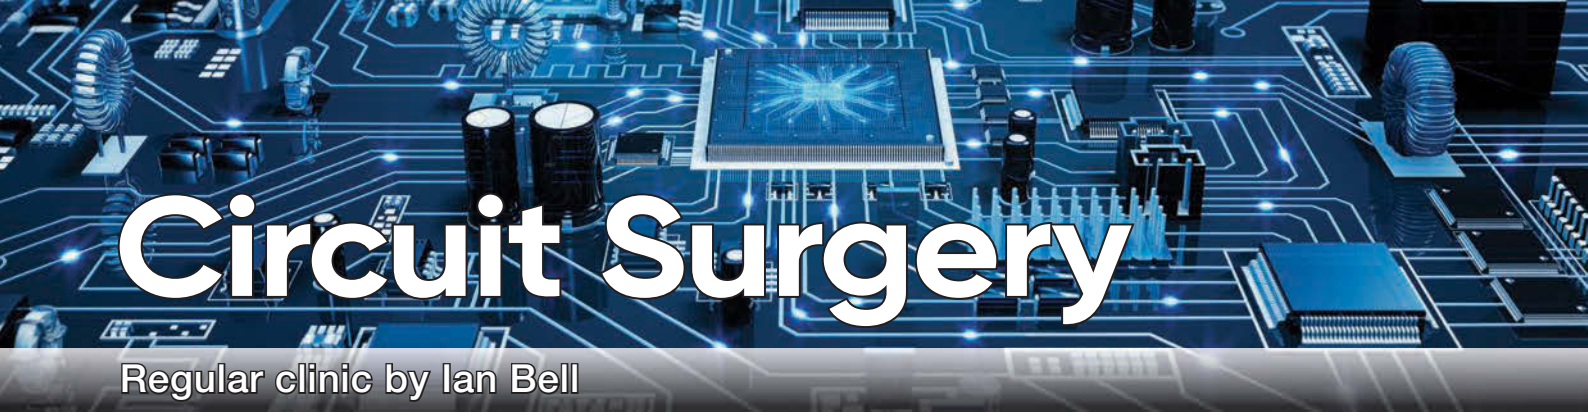

Circuit Surgery

Regular clinic by Ian Bell

A transistor-level model of the 555 IC in LTspice

Last month, we looked at making resistance measurements using a 555 timer IC in the context of plant music devices. The oscillation frequency of the 555 astable (oscillator) depends on the resistance of the plant to which it is connected, which takes the place of one of the timing resistors in the standard astable circuit.

The frequency is easily measured using a microcontroller by triggering interrupts from the 555's output.

The music device measures fluctuations in electrical resistance of the plants (or anything conductive) and translates the results into either MIDI (Music Instrument Digital Interface) commands or directly into synthesised music audio, via either a built-in speaker or a phone app. This process is often called (bio)sonification.

The 555 chip dates from the early 1970s – it had its 55th anniversary this year. It has been described as one of the most popular ICs of all time, and last month we briefly considered its history and importance.

Its usage has declined in recent decades because many typical 555 roles are now easily implemented with low-cost microcontrollers. Despite this, it is still produced and used in large numbers.

This month, we will look into further aspects of the 555 chip, using it for resistance measurement and seeing how we can model and simulate the 555 in LTspice using a more realistic model than the one we used last month. The discussion on using the 555 model has wide relevance for other devices.

Water conductivity

When investigating plant music circuits, I came across several examples of a similar 555 circuit used for measuring the conductivity of water (or other conductive liquids, or soil), usually in environmental monitoring contexts. A common configuration is shown in Fig.1. This differs from the plant music circuit in that it uses a different 555 astable configuration and employs capacitive coupling to the probes (C_{isol1} and C_{isol2} in Fig.1).

Conductivity (σ, Greek letter sigma) is a property of a material that defines how well it conducts electricity. The reciprocal quality is resistivity (ρ, Greek letter rho), for which $\sigma = 1 \div \rho$.

These can be defined with reference to Fig.2, which shows a length L of uniform material of conductivity σ (resistivity ρ). The material has a uniform cross-sectional area A and a voltage V is applied across its length via ideal electrodes at the two ends, as shown in Fig.2. The material could be a solid or a fluid held in an insulating container with cross-sectional area A.

Applying the voltage results in a current flow, I. From Ohm's law, the resistance (R) of the cylinder of material in Fig.2 is $R = V \div I$. The resistance of a specific cylinder will be proportional to its length and inversely proportional to its cross-sectional area. Increasing the length is like adding series resistance in a circuit, whereas increasing the area is like adding parallel resistance.

The constant of proportionality between resistance, length and area is resistivity. We can also write the relationship in terms of conductance, $G = 1 \div R$, with conductivity as the constant. We have:

$$R = \rho\frac{L}{A} \quad \text{and} \quad G = \sigma\frac{A}{L}$$

Of course, we can also find resistance in terms of conductivity and conductance using resistivity, where convenient. Conductivity is measured in siemens per metre (S/m) and resistivity is measured in ohm-metres (Ω·m).

The circuit in Fig.1 directly measures resistance, not conductivity or resistivity. If we just randomly hold two probes in some water, we can measure resistance, but we cannot find the conductivity because we do not have a known physical relationship between the probes.

We need to build a pair of electrodes with a fixed physical structure and a defined relationship to the water; for

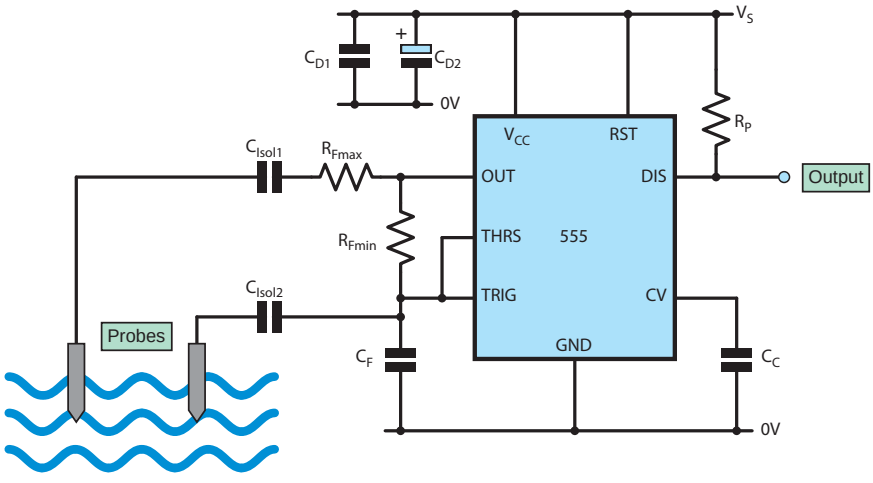

Fig.1: making water conductivity measurements using a 555 IC.

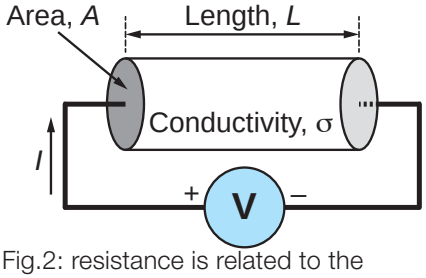

Fig.2: resistance is related to the conductivity of the material, its cross-sectional area and length.

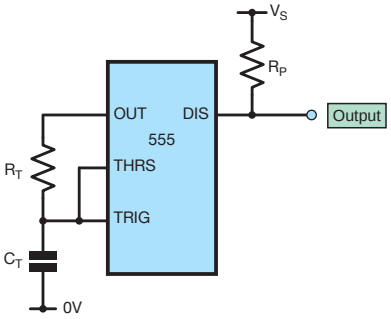

Fig.3: an alternative 555 astable circuit configuration.

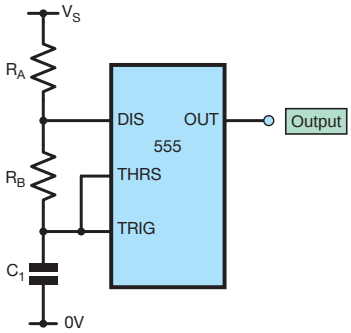

Fig.4: the standard 555 astable configuration.

example, electrodes in a tube through which the water flows, or fixed in a probe designed to dip into bulk water. With the physical structure fixed, the system can be calibrated to measure conductivity.

Conductivity measurement of water is commonly used in environmental science and pollution monitoring to provide information on water quality (for freshwater) or salinity (for seawater).

Pure water is a poor conductor (low conductivity), but dissolved salts increase its conductivity. Salts are substances such as sodium chloride and calcium carbonate that ionise when dissolved. The ions provide the charge carriers that increase the conductivity.

Natural fresh waters (rivers, lakes etc) contain naturally occurring salts, depending on the local geology, which is typically fairly constant under normal conditions. Pollution discharges and runoff, saltwater intrusion into freshwater, extreme rainfall, and other such events often result in large and rapid changes in conductivity.

These can have a significant impact on ecosystems and human health, but can be effectively monitored using conductivity measurements.

Conductivity instruments used in professional environmental research and monitoring are likely to be significantly more advanced than the circuit in Fig.1, but it can be used in less-critical applications such as citizen science projects, education and gardening/horticulture (eg, soil moisture measurement).

Fig.5: the simplified internal arrangement of the 555 IC.

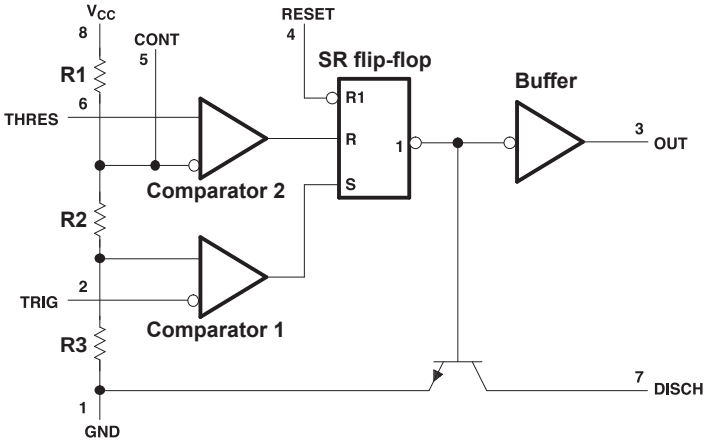

A different astable configuration

The circuit in Fig.1 is a 555 astable multivibrator, but it is different from the most commonly used configuration, which we discussed last month. Fig.3 shows a simplified version of the circuit in Fig.1 that leaves out 555 pins not directly involved in the oscillation timing and output. It also just shows R_T as the timing resistor, rather than the capacitively isolated probe circuit from Fig.1.

Fig.4 shows the standard 555 astable circuit from last month, drawn in the same format as Fig.3 for comparison.

The alternative astable circuit in Fig.3 does not use the DIS pin to discharge the timing capacitor during the output-low part of the oscillation cycle, as is done in the standard circuit. Instead, it drives the timing circuit directly from the 555's output.

When OUT is high (at V_s), the timing capacitor C_T charges via R_T until it reaches $\frac{2}{3}V_s$, at which point the comparator connected to the THRS input (see Fig.5) switches, resetting the flip-flop, and causing OUT to go low (0V). C_T will then discharge via R_T.

C_T discharges via R_T until it reaches $\frac{1}{3}V_s$, where the comparator connected to the TRIG input (see Fig.5) will switch, setting the flip-flop, and the sending the OUT pin high. C_T starts charging again, and the oscillation cycle repeats.

This behaviour is similar to the standard astable in that the timing capacitor voltage ranges between $\frac{1}{3}V_s$ and $\frac{2}{3}V_s$, but there is only one timing resistor involved, so the high and low output times are equal (under ideal circumstances).

This is an advantage of this circuit, along with its use of fewer timing components.

We have to add the pull-up resistor (R_P), although this may not be required when interfacing to a microcontroller GPIO input with an internal pull-up.

The DIS pin's open-collector/open-drain output is used to provide the circuit's output using a pull-up resistor (R_P) to obtain a high (logic 1) output when the internal transistor is off. This avoids the OUT pin being connected to external circuitry, which might influence the timing behaviour if the external circuity's impedance were not sufficiently high.

The OUT and DIS pins are both driven from the same signal inside the 555 (the Q-bar output of its RS flip-flop – see Fig.5), so they essentially produce the same waveform with R_P in place.

Frequency and duty cycle

Each charge/discharge cycle takes $0.693R_TC_T$, so the period of oscillation is twice this ($1.386R_TC_T$) and the frequency (f) is approximately:

$$f = \frac{1}{1.4\,R_T\,C_T}$$

Duty cycle is the proportion of the cycle period that the pulse is high. A pulse train with a 50% duty cycle is an ideal square wave, which is high for exactly half the period. Lower duty cycles have proportionally shorter pulses, while higher ones have longer pulses. As already noted, the circuit in Fig.3 produces a 50% duty cycle.

For the standard astable configuration (Fig.4), the total cycle time is $T=0.693(R_A+2R_B)C_1$ and the time the output is high is $T_{ON} = 0.693(R_A+R_B)C_1$, so the duty cycle is:

$$\text{Duty cycle} = \frac{T_{ON}}{T} = \frac{R_A+R_B}{R_A+2\,R_B}$$

If we make R_A relatively very small, and R_B very large, the duty cycle will be close to 50% ($\gg R_B \div 2R_B$), but it cannot be smaller than this, or exactly equal to 50%.

The square-wave output is a useful property of the circuit in Fig.3. Although it is not essential for the conductivity measurement system, it may be important in other applications.

Simulation

Fig.6 shows an LTspice schematic of the alternative 555 astable circuit from Fig.3. Using the formula above (with $1.386R_TC_T$, as this is an idealised 555) we expect a frequency close to 1kHz. This is confirmed by the simulation results shown in Fig.7. We can also see that the output has a 50% duty cycle, as discussed above.

The signals on the DIS and OUT pins are equal, as expected, thanks to the pull-up resistor (**RP**) on the DIS pin.

On the lower pane, the voltage on the timing capacitor, **V(TrigThresh)**, is shown together with the comparator switching thresholds: $\frac{2}{3}V_s$ = **V(CV)** and $\frac{1}{3}V_s$, which is plotted using waveform maths (**V(CV)/3**), as we did last month. This shows that the capacitor charges and discharges between $\frac{1}{3}V_s$ and $\frac{2}{3}V_s$ as it does in the standard astable circuit.

555 models

The built-in model of the NE555 in LTspice is idealised. This is useful for working with the fundamental behaviour of the device to understand its operation, but it will not show non-ideal performance that may undermine some designs. Nor will it show whether mitigations for any possible problems are effective.

Fortunately, there is a full model of the NE555 included as part of the LTspice download ("for informational/educational purposes only"). The default location on Windows for LTspice v24 is: **C:\Users\<username>\AppData\Local\LTspice\examples\Educational\NE555.asc** (replace **<username>** with your user name).

On other operating systems or with other versions, it may be in a different location.

Fig.8 shows the left-hand side of the schematic, which includes most of the 555's pins. The output is on the right-hand side. This schematic is not just the NE555's transistor-level circuit; it also includes external components for a standard astable running from a 15V supply (V1), with the reset pulsed on and off using the **V3** source.

The schematic is configured to run a 30ms transient simulation, during which the reset pulses on and off three times. Fig.9 shows one of the three bursts of output pulses controlled by the reset. **V(3)** (yellow trace) is the pin 3 (output) voltage, while **v(2)** (red trace) is the pin 2 (trigger) voltage, which is the voltage at the timing capacitor.

A couple of things are worth noting. The first pulse in each batch is longer than the others because the timing

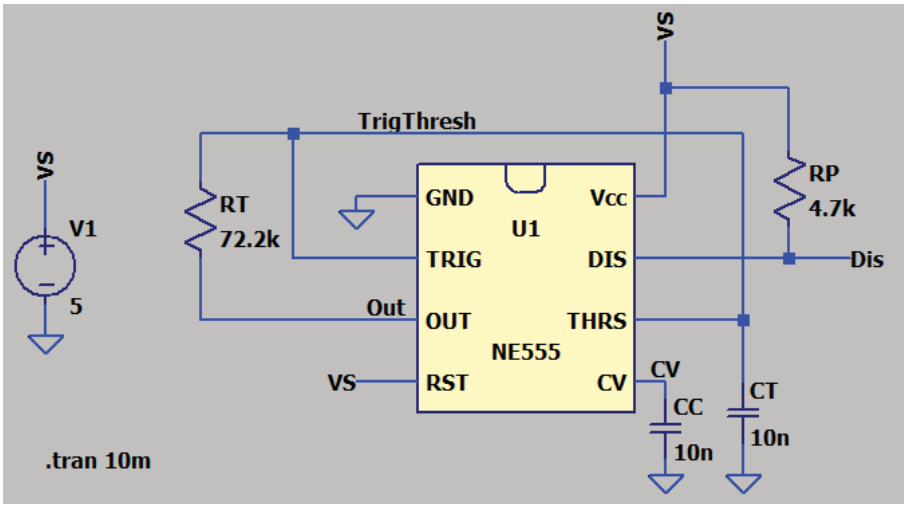

Fig.6: an LTspice schematic of the alternative 555 astable configuration in Fig.3.

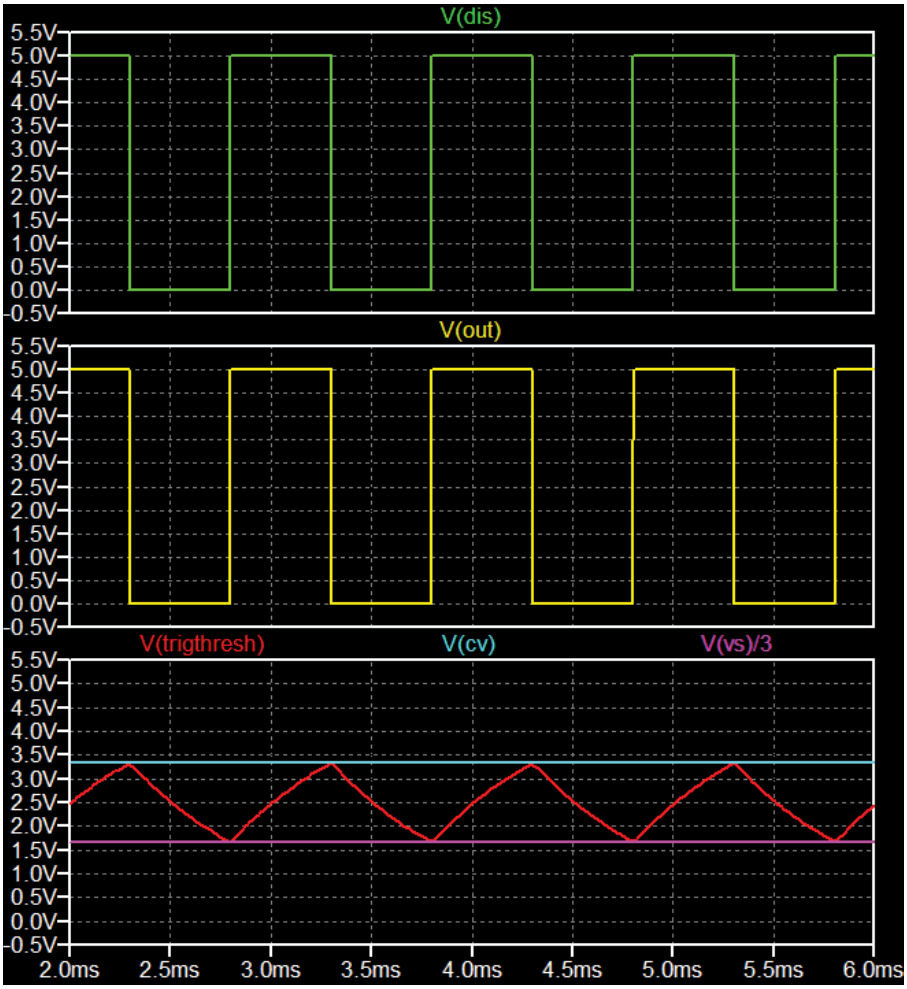

Fig.7: the simulation results from the circuit in Fig.6.

capacitor is starting from 0V rather than one-third of the supply voltage. Secondly, the output does not go fully to the supply voltage; it reaches around 14.6V rather than 15V. This is not unexpected; many ICs are not able to drive their outputs to equal the supply voltage.

The NE555 datasheet states a worst-case output high voltage of 12.75V on a 15V supply, but that is with a 100mA output current. The output in the example circuit is not connected to anything, so the voltage is not always as high as this.

Hierarchy and subcircuits

It would be useful to be able to use the NE555 transistor-level circuit as a component in other simulations; for example, so that we can investigate the effect of its non-ideal behaviour on the alternative astable configuration. To do this, we need to convert it for use as a hierarchical schematic or library subcircuit file.

Before looking at the details of how to do this, it is useful to define what these two options are. Both are ways in which devices other than LTspice's built-in fundamental components such

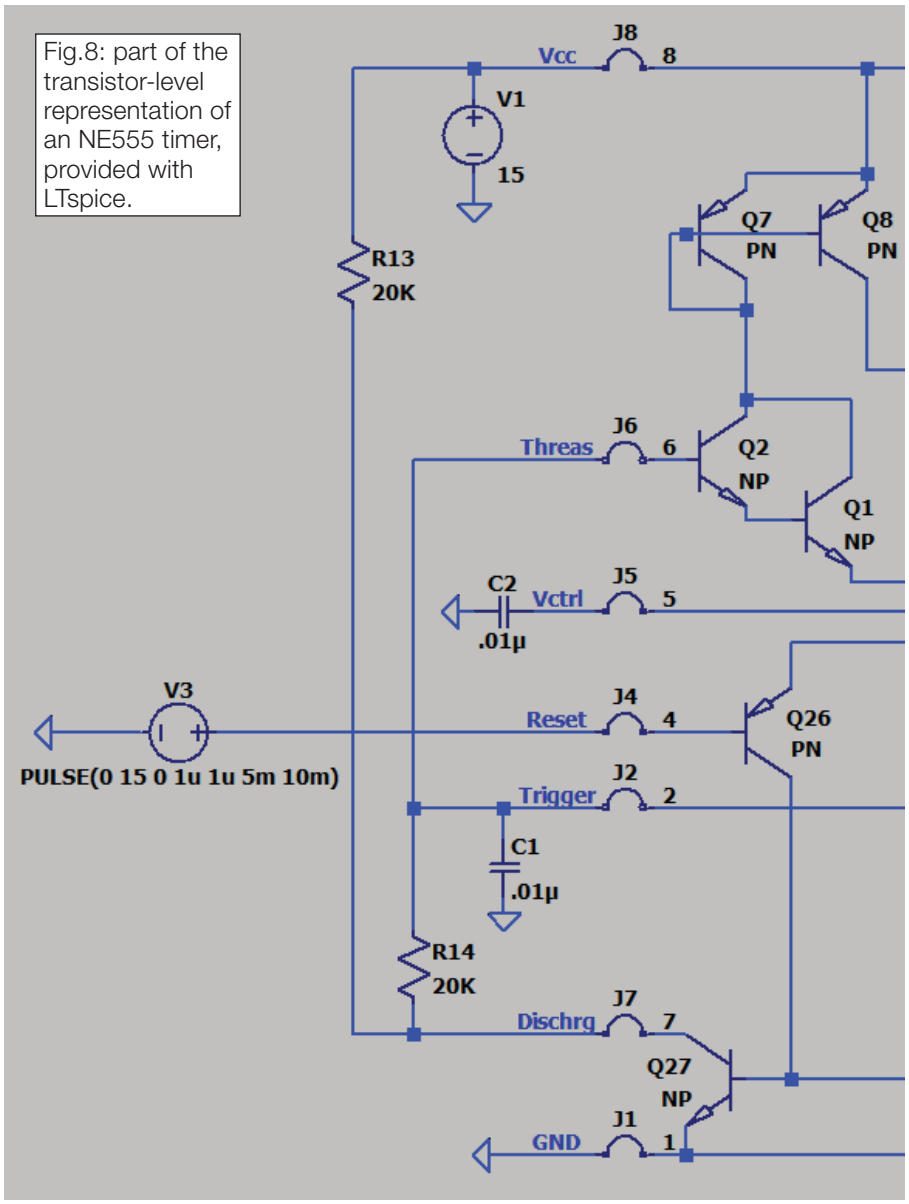

Fig.8: part of the transistor-level representation of an NE555 timer, provided with LTspice.

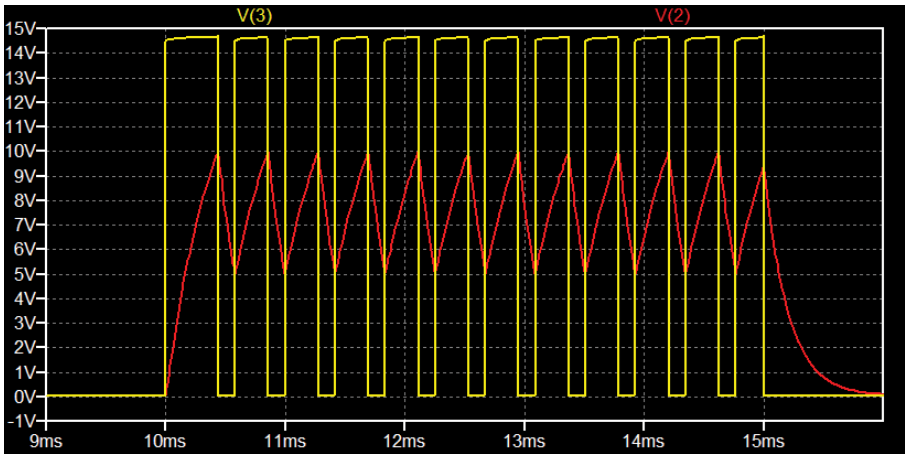

Fig.9: the simulation results from the circuit in Fig.8.

as resistors, capacitors and diodes can be included in simulations.

Hierarchical schematics are where a symbol on one schematic links to another LTspice schematic file. This can happen over multiple levels – a schematic referenced in one file can in turn include a symbol that references another schematic.

The same lower-level schematic can be referenced multiple times from different symbol instances in a schematic; LTspice will give individual identities to the components in the different copies of the referenced schematic. LTspice combines all the schematics into a single circuit for simulation.

LTspice does not use the schematic drawings directly as input to the simulation engine. It converts the schematic drawing to a text-based description of the electrical circuit called a netlist, which is passed to the simulator. You can view the netlist for a schematic using the menu option **View → SPICE Netlist**.

A netlist does not contain any information about the graphical structure or look of the schematic – that is not needed by the simulator. You do not need a schematic to create a netlist or simulate a circuit. SPICE simulators predate modern graphical user interfaces – in the early days, all circuits were input as text netlists.

Complex devices such as ICs can be described using LTspice netlists placed in text files (usually with .sub or .lib extensions). This is how the behaviour of the thousands of Analog Devices products (such as op amps) and other complex components is defined in the library provided with LTspice.

A raw schematic or netlist of a circuit is not quite sufficient to enable its use in a larger circuit – it is necessary to define which nets (wires) in the circuit are the external connections (inputs, outputs and supplies). On a schematic intended for hierarchical use, this is achieved by placing 'ports'.

For a netlist, the circuit description is placed between `.subckt` and `.ends` directives and the `.subckt` command declares the subcircuit name and lists its external connections. Hierarchical schematics are automatically netlisted using `.subckt` directives.

Both approaches allow a symbol to be linked to the circuit file (schematic or subcircuit text). This symbol can then be used to represent the device on schematics in the usual way. Symbol creation is not as onerous as it may sound because symbols can be generated automatically or existing symbols can be copied and edited.

Subcircuits instances (represented by their symbols on a schematic) are netlisted like other components in an LTspice circuit. Just as resistor names always start with an 'R' and capacitors with a 'C', subcircuit have 'X' as the first letter of their name. This is handled automatically by LTspice.

Schematic conversion

Of the two options, using a hierarchical schematic is more straightforward and may therefore be best for one-off use. A library file is more flexible and is probably the best method if the model is to be used multiple times in different designs.

The first step is to copy the file from the LTspice AppData folder noted above to a folder in your user file space. Given that the filename is the same

as the idealised model (NE555), it is a good idea to rename it; for example, to **NE555_BJT_sch.asc** (indicating it's the bipolar transistor schematic version).

After that, we need to edit it to remove all the external components, voltage sources and simulation commands.

The jumpers (**J1**, **J2** etc) seen in Fig.8 are a special schematic symbol (in the Misc section of the library). Jumpers do not create a circuit element (component) in the simulation, but are used to inform LTspice that there are two different names for the same electrically identical node.

Looking at Fig.9, at first it might seem that the pins have been labelled with both their pin numbers and their names (eg, **Threas** is pin 6), using the jumpers to facilitate the multiple naming. However, the names are just text comments on the NE555 schematic.

I assume the jumpers are used to mark the boundary between the internal and external circuitry. Certainly, this is helpful when editing the schematic to remove the external components. The jumpers themselves do not need to be removed.

Hierarchical schematic

To use one schematic in another, the lower-level schematic (our NE555 circuit in this case) must have its input and output ports defined. This is done when labelling nets by defining the port type (right-click a wire and select **Label Net**).

You need to select **Input**, **Output** or **Bi-Direct** in the **Port Type** dropdown. This will provide a port symbol to place on the schematic instead of the default, which is just the net label text. You need to set up all external connections to the circuit as ports, including power supplies and ground.

Fig.10 shows the NE555.asc schematic edited to add the port symbols; the output pin (OUT) is shown for completeness although it is on the other side of the diagram. The external components and simulation command have been removed.

The original pin-number net names have been removed to prevent possible confusion. The pins have been added with the same names as the pins on the built-in ideal 555 model – this will make using a copy of its symbol easier.

What port type to use is a possible source of confusion. For example, power supplies and ground are typically thought of as a separate type of connection, not as inputs or outputs (they are configured as inputs on the schematic in Fig.10).

Some connections do not fit neatly into a simple input or output

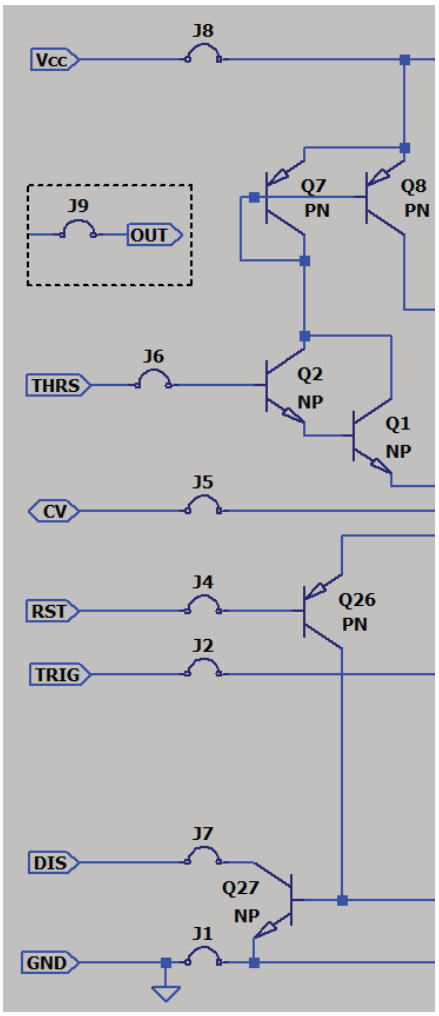

Fig.10: part of the 555 schematic from Fig.8 modified for use as a subcircuit.

classification. For example, the 555's control-voltage (CV) pin is connected to an internally generated reference voltage, so it could be used as an output, but an external voltage may also be applied to it to alter the comparator thresholds, making it act as an input. It therefore is probably best described as bidirectoinal.

By contrast, the DIS pin is an open-collector output controlled by the internal latch, even though current normally flows into the pin when it discharges the timing capacitor.

Fortunately, it does not really matter what port type you use. LTspice only uses port type information for generating symbols and determining overall connectivity for hierarchical schematics. It does not use port types during the actual simulation process.

In the simulation, all points in a hierarchical circuit connected together via ports are treated as electrical nodes, whatever the port types they happen to be assigned. In the simulation, there is no implication of signal direction or electrical properties of the circuit implied by the port type.

Once the schematic has its ports defined, a symbol can be generated for it

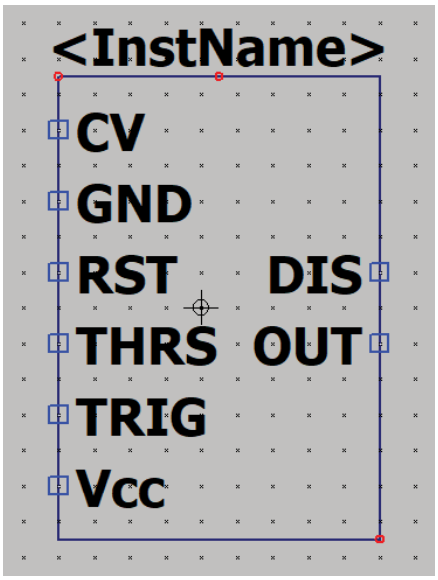

Fig.11: the automatically generated LTspice symbol from the 555 circuit.

using the menu item **Hierarchy → Open this Sheet's Symbol** and responding to the prompt to generate a new symbol. If there are no ports on the schematic, you will receive an error message. The symbol will be created in a LTspice symbol file (.asy extension) with the same main name as the schematic.

Fig.11 shows the symbol automatically generated from Fig.10. If you change the port configuration and want to generate a new symbol, you need to delete the existing symbol file. Otherwise, the existing one will simply be opened rather than a symbol creation being offered.

Symbol copying

The symbol in Fig.11 is perfectly usable but may not be ideal. The symbol can be edited to change the position of the port and its overall shape. However, if we want the symbol to match the existing one for the ideal 555, it is easier to use a copy of it instead.

This means we get a symbol for the NE555 schematic that has the same shape and pin locations as the built-in ideal 555, which makes it straightforward to create variants of schematics using the 555 with different models of the device for comparison purposes.

The existing 555 symbol is at the location **C:\Users\<username>\AppData\Local\LTspice\lib\sym\Misc\NE555.asy** or similar for different versions or operating systems. To use this symbol, copy the file to the same folder as the NE555 schematic and rename it so that it has the same name as the schematic file, with a different extension (eg, **NE555_BJT_sch.asc** and **NE555_BJT_sch.asy**).

Open the symbol file in LTspice (see Fig.12). It is configured for use with a library file subcircuit. We need to

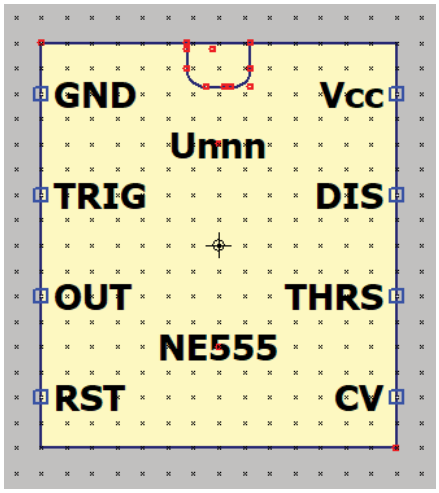

Fig.12: the NE555 symbol included with LTspice.

change the symbol type and remove the attributes that are used with text-based subcircuits.

One of these attributes is the name "NE555" on the symbol. Delete this text first. The delete key may not activate the delete function as it does in the schematic editor, so click the delete icon ⓧ on the toolbar or use **Edit → Delete**.

Replace the deleted device name with a new text comment and different name via the t icon on the toolbar or **Draw → Text**. A text comment will not be interpreted as an attribute.

To deal with the other attributes, use **Edit → Attributes → Edit Attributes**, which opens the attribute editor. Change the "Symbol Type" dropdown from Cell to Block and remove all the attribute values (double-click the text in the value column to select it and then click the delete key) – see Fig.13.

The Block symbol type is used for hierarchical schematics, whereas the Cell type is for library subcircuit text files. If the attributes are not blank, LTspice may produce errors when a symbol links to a hierarchical schematic.

Click OK in the attribute editor and save the symbol. The symbol's yellow background will disappear, and the **Unnn** instance name will be replaced

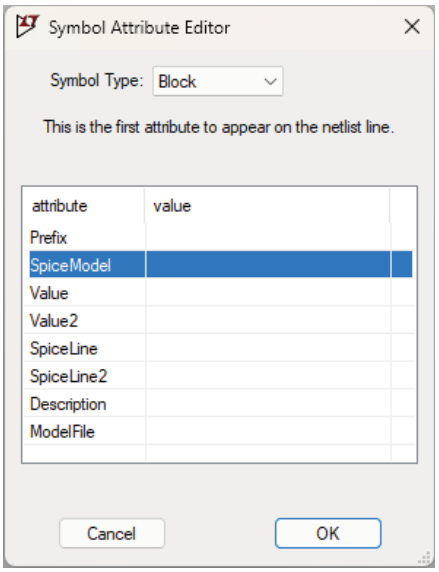

Fig.13: the symbol attribute editor for a hierarchical schematic.

by **<InstName>**. It is the Prefix: X attribute, required for library subcircuits, that causes a yellow symbol background (assuming there is a rectangle in the symbol).

The symbol can now be used in a schematic in the same folder as the NE555 schematic and symbols files. To place the symbol, you use the component tool from the toolbar in the usual way. To find the symbol you need to change the "Top Directory:" dropdown in the component dialog to the local folder instead of the built-in LTspice library location.

Fig.14 shows an example of the new symbol in use. Compared with Fig.6, the 555 symbol has a name starting with an "X" rather than a "U". These are just the default names on the schematic drawing – LTspice uses U for text subcircuits. The actual netlist adds the X prefix to both hierarchical schematics and subcircuits (eg, to get XX1 or XU1).

Library subcircuits

To create a library subcircuit from the NE555 transistor-level schematic, we start by creating a netlist text file for the circuit. To do this, modify the original circuit as described above to remove the external components (if you have not already done so).

With the modified circuit (as partly shown in Fig.10) open in LTspice, select **View → SPICE Netlist** menu option. In the text area, right-click and select "Edit as Independent Netlist". Save the file via the dialog that appears, using a different main name from the schematic (eg, **NE555_BJT** rather than **NE555_BJT_sch**) and with the extension **.lib** rather than the default **.cir**.

The netlist will open the LTspice text editor as shown in Fig.15. This is a schematic netlist, but we need a subcircuit definition, which requires some edits to the file. The syntax of a subcircuit definition is:

```
.subckt <subcircuit_name> pin1 [pin2 ...]
...
.ends [subcircuit_name]
```

pin1, **pin2** etc are the names of external connections as used in the netlist. The order of the pins is important and must match the order defined for the symbol used for the subcircuit. The **.ends** directive must be the last line of the subcircuit and may optionally repeat the name. Library files may contain multiple subcircuits, so this is useful for keeping track of things.

To edit the NE555 file, we need to add the line:

```
.subckt NE555_BJT 0 TRIG OUT RST CV THRS DIS Vcc
```

…before the line with the first transistor (**Q1**). The pin names and order here match those used in the existing NE555 symbol, to facilitate its reuse. We also need to delete the **.backanno** directive as it is not needed here, and to change the **.end** directive to **.ends**. Having made these changes, save the file.

As with the hierarchical schematic, we need a symbol for our subcircuit

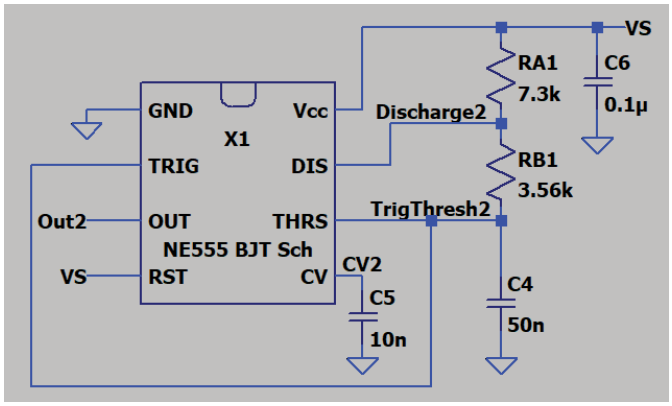

Fig.14: an NE555 astable circuit using the transistor-level sub-circuit in a hierarchical schematic.

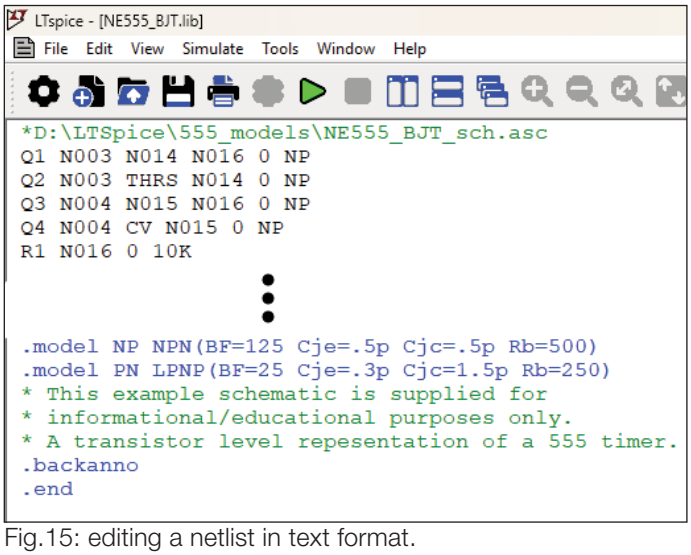
LTspice - [NE555_BJT.lib]
File Edit View Simulate Tools Window Help
*D:\LTSpice\555_models\NE555_BJT_sch.asc
Q1 N003 N014 N016 0 NP
Q2 N003 THRS N014 0 NP
Q3 N004 N015 N016 0 NP
Q4 N004 CV N015 0 NP
R1 N016 0 10K
.model NP NPN(BF=125 Cje=.5p Cjc=.5p Rb=500)
.model PN LPNP(BF=25 Cje=.3p Cjc=1.5p Rb=250)
* This example schematic is supplied for
* informational/educational purposes only.
* A transistor level repesentation of a 555 timer.
.backanno
.end

Fig.15: editing a netlist in text format.

model. Again, this can be automatically generated. To do so, right-click the **.subckt** line in the file and select **Create Symbol** from the menu. As before, it may be better to adapt the existing symbol, as the generated one will not have the same shape.

As above, copy the symbol file from the LTspice library to the same folder as the subcircuit library file and rename it so that it matches the main part of the library filename (eg, **NE555_BJT.lib** and **NE555_BJT.asy**).

Open the symbol file (see Fig.12), right-click on the "NE555" and change it to make the name match the subcircuit name in the **.subckt** directive within the .lib file (eg, **NE555_BJT** if it is the same as the example above).

Open the attribute editor (**Edit → Attributes → Edit Attributes**) and adjust the values as shown in Fig.16. The "Symbol Type" and prefix should already be "Cell" and "X", respectively. "Value" should have been updated by the previous edit of the text on the symbol. "SpiceModel" can be changed to "NE555_BJT" or deleted. "Value2" can be deleted.

The "Description" should be changed to accurately reflect this circuit and "Modelfile" needs to be set to the filename of the library file.

Either or both of the "Value" and "SpiceModel" values can be the subcircuit name in the library file. If both are blank, LTspice will be unable to find the model. The "Modelfile" attribute can be just the filename if the library file is in the same folder as the schematic in which it is used; otherwise, the full path is required.

Grounding

The first pin on the symbol is labelled **GND**, but the subcircuit netlist uses the standard net name of **0** for the ground (reference) net. The **.subckt** directive above used **0** rather than **GND** to ensure the pin name matches the netlist. If **GND** were used in the **.subckt** directive, the pin would not connect to anything, as the net name **GND** does not occur in the transistor-level NE555 netlist.

The net name **0** does not match the symbol pin name of **GND**, but that does not matter because it is the order in which the pins are listed in the **.subckt** directive that determines which symbol pin they are mapped to.

To view or configure the symbol pin to netlist mapping, open the symbol for editing and right-click the pin name text on the symbol. This will open the Pin/Port Properties dialog (Fig.17). The Label box is used to set the pin name displayed on the symbol. The "Netlist Order" spinbox sets which of the pins

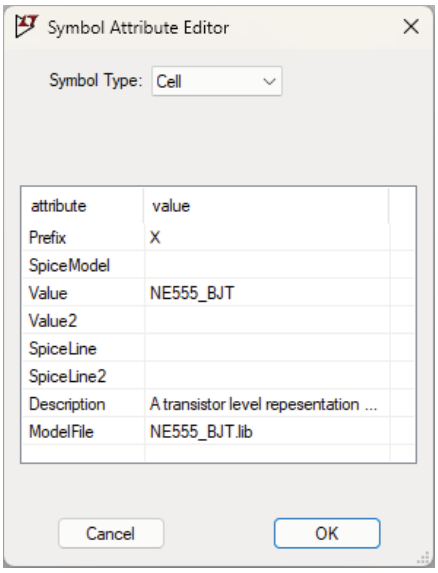

Fig.16: using the symbol attribute editor on the NE555 subcircuit.

listed in the **.subckt** directive maps to this pin.

Fig.18 shows a version of the circuit in Fig.6 with a 1Ω resistor inserted between the 555's ground pin and the circuit ground (**RM_1** for ground monitor). This means that the 555's ground is not connected to LTspice's global ground 0V reference point (node **0**) by the wiring of this schematic.

Although IC ground pins are often connected to a circuit's ground, this is not always the case, and model should behave appropriately in this situation.

The full schematic, of which Fig.18 is a part, contains three copies of the circuit shown in Fig.18 with three different models of the 555. All other connections and components are the same. The components and nets are labelled with ***_n*** after their name, where ***n*** is 1, 2 or 3 for 555 devices U1, U2 and U3 respectively. The 555 models are:

- U1: The built-in idealised model.
- U2: The BJT schematic model with its first pin named **0**.
- U3: The BJT schematic model with its first pin named **GND**.

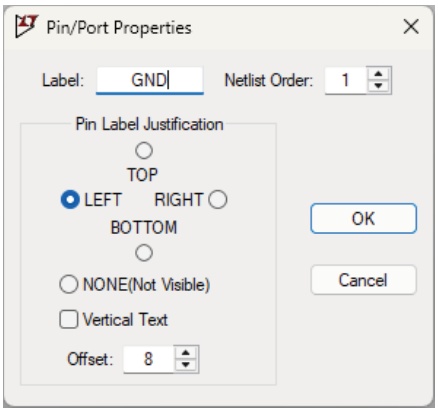

Fig.17: the symbol Pin/Port Properties dialog.

The pin names refer to the discussion above on the name used for the 555's ground pin in the **.subckt** directive.

Fig.19 shows the simulation results from Fig.18 – these are plots of the voltage at the ground pins of the 555, which will be proportional to the current through the pin. U1 (green trace) and U2 (magenta trace) show a pulsing waveform, which is to be expected as the 555 switches its output.

The waveform for U3 (red trace) shows no activity, although its output waveform exactly matches that for U2 (not shown). U3 is working correctly with its ground connected via the global node '**0**' inside the subcircuit, not via its pins. The pin is open-circuit due to the use of the **GND** pin name, as discussed above. This is confirmed by a message in the SPICE log:

Ignoring empty pin current: Ix(u3:gnd)

The log is worth checking (**View → SPICE Output Log**) when setting up subcircuits, as messages there may alert you to problems that are not obvious from the simulation results.

The fact that node '**0**' is a global ground may make you think something strange has happened in this simulation. If the ground in the 555 subcircuit uses node '**0**', why is it not connected to the global ground inside the subcircuit, effectively shorting out the monitor resistor? In this case, all

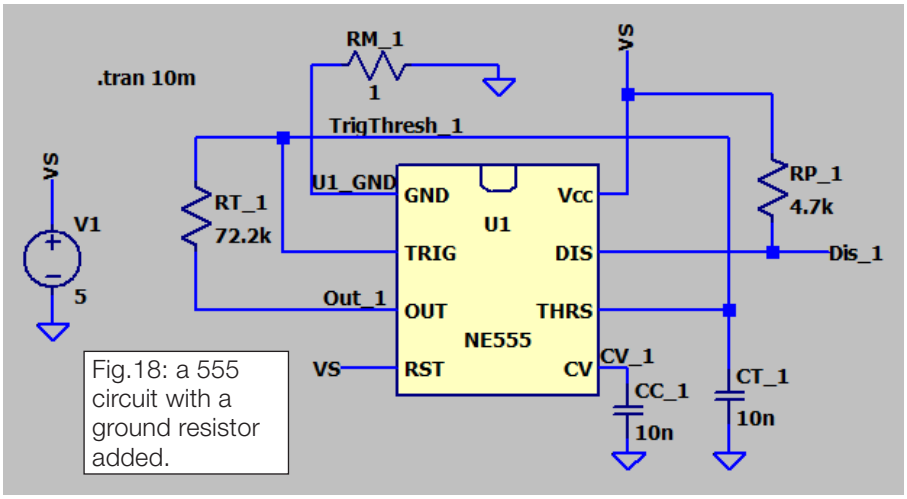

Fig.18: a 555 circuit with a ground resistor added.

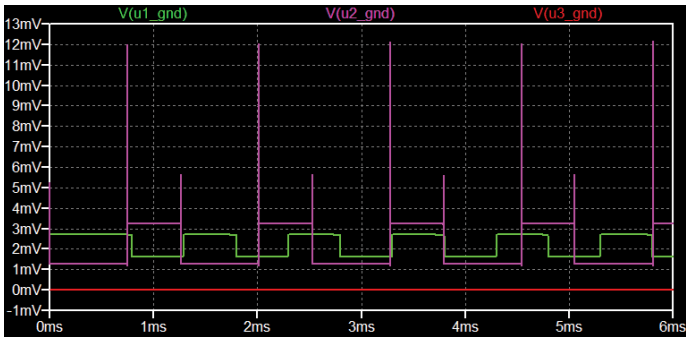

Fig.19: the simulation results for the circuit in Fig.18.

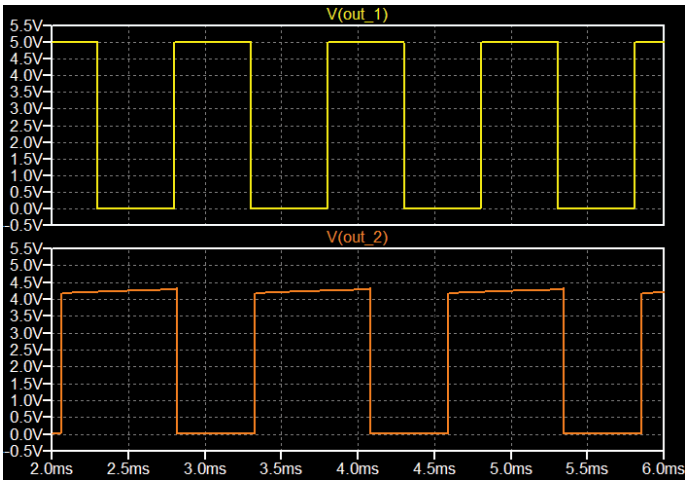

Fig.20: the output waveforms for Fig.18 with the ideal and transistor-level NE555 models.

the signals in Fig.19 should be constant 0V. The LTspice help indicates that this should happen:

*The .global command allows you to declare that certain nodes mentioned in subcircuits are not local to subcircuit but are absolute global nodes. Note that global circuit common is node '**0**' and that a .global statement is not required.*

It may not be documented, but LTspice treats node '**0**' used as a subcircuit pin name as a special case and does not connect it to global ground. Instead, it behaves the same way as other subcircuit pins and takes the name of the net in the higher-level circuit to which it is connected.

We can view the final flattened netlist of a circuit by selecting the **View → Netlist** menu option, then right-clicking the text and selecting "Generate Expended Listing". Save the file when prompted, and it will open as a text file in LTspice. The bottom resistor of the three-resistor potential divider (R1,R2,R3 in Fig.5) is **R5** in the NE555_BJT subcircuit netlist. It is connected to nets **N025** and **0**:

```
R5 N025 0 5K
```

In the expanded netlist for Fig.18, this resistor in **U2** is not connected to **0**, but to net **u2_gnd**, which is where the 555's ground is connected in the schematic:

```
r:u2:5 u2:n025 u2_gnd 5k
```

For **U3**, the resistor is grounded, as in the original subcircuit:

```
r:u3:5 u3:n025 0 5k
```

This fits with the results shown in Fig.19.

This behaviour by LTspice is useful. It means that we can convert a standard schematic for a device with node '0' as ground into a library subcircuit without having to rename all the ground connections to prevent the device's ground pin from being shorted to ground when it is actually connected elsewhere.

Circuit performance

Fig.20 shows the output waveforms on the **OUT** pins of **U1** and **U2** for the circuit in Fig.18, with the ground monitor resistors changed to 0Ω to prevent any impact on the behaviour of the circuit. We can see a significant difference between the ideal 555 model (top pane) and the transistor-level model (bottom pane).

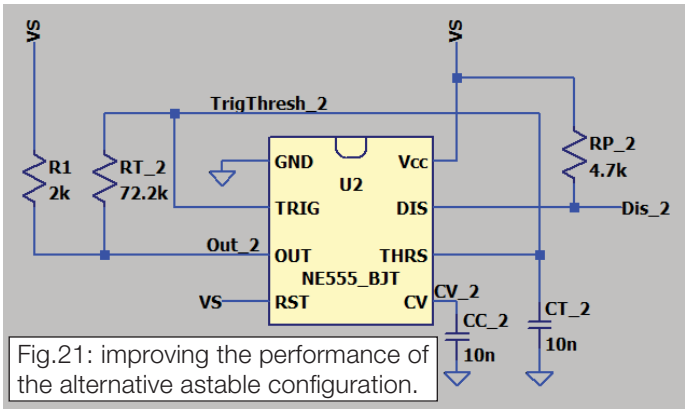

Fig.21: improving the performance of the alternative astable configuration.

For the transistor-level model, the frequency is lower and the duty cycle is not close to 50%. The low output voltage goes close to 0V, but the high level is significantly lower than the 5V supply.

The fact that a 555's output does not reach the full supply voltage is the key reason for the change in frequency and duty cycle. The timing capacitor is charged from the output pin, so it will take longer to reach $\frac{2}{3}V_s$ when the output is high because it is heading towards a lower final voltage.

The low output voltage is close to zero, so the discharge time will be less severely affected. The high output becomes longer, while the low output is similar to the ideal case, so the duty cycle is above 50% and the frequency is lower.

This problem can be overcome by connecting a resistor from **OUT** to the supply, as shown in Fig.21 (**R1**). This pulls the output pin up to the supply voltage when the output is high. Its value needs to be much smaller than the timing resistor, but if it is too small, it will raise the voltage level during the low output. The results in Fig.22 show that the circuit in Fig.21 provides a significant improvement.

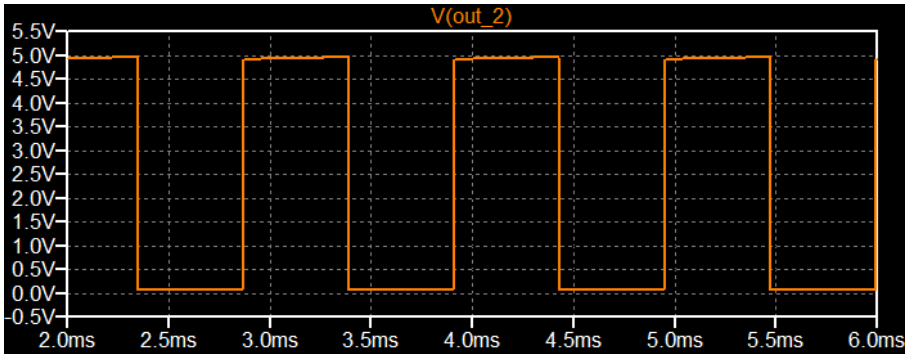

Fig.22: the simulation results for the circuit in Fig.20.

Conclusion

We have looked at simulations of a 555 circuit using an idealised model and one based on a transistor-level schematic. The latter provides more accurate results, but took some effort to set up.

Some aspects of creating library models in LTspice, such as the behaviour of ground connections, may be a little non-obvious at first. However, like many things, after some experience and practice, creating subcircuit models from existing schematics is straightforward. PE

Real Live Robots

My hands-on experience with robots has so far been confined to the Lidl/Parkside robot lawnmower that I reviewed in the July 2025 issue. I've seen plenty of movie clips of humanoid or animaloid bots playing games and dancing, but never had the opportunity to watch them at work live in a real-world situation.

So, I jumped at the chance to attend a conference and demonstration on "Advancing Commercial Deployment of Humanoid Robotics in Europe" held recently at the new(ish) events centre, The Magazine, near the O2 dome in Greenwich.

The event was staged by Chinese company Agibot and I left with two very solid convictions. Although some robot technology may be based on open standards, the Chinese look likely to lead the world in practical implementation, because there is clearly a money tree growing in China to make robots a future way of life.

Secondly, young people hoping for a secure future will need either top-level creative brain skills, such as software design and AI development, or practical trade skills in fields like plumbing, electrical maintenance or mechanical repair.

While robots will be able to master repetitive tasks and basic customer-facing jobs, it is hard to envisage robots under sinks, up ladders or in awkward cupboards fixing leaky washers, failing fuse boxes or air-locked central heating.

In the future, new homes may well be designed with standard pipework, cables and componentry that robots can work with. But that will take decades and there will always be a lot of people who prefer to live in quaint old houses with retro equipment hidden in hard-to-access nooks and crannies.

Agibot of Shanghai says that although the company is only three years old, it has already produced 15,000 working bots. Its technology is now moving from the evaluation to deployment stage, on the road to robot autonomy, with 2030 seen as the year when robots 'universalise' and become a familiar sight.

By then, travellers will be used to seeing a humanoid bot working at the check-in desk of a hotel, understanding over 100 languages and communicating in thirty. Bots will have replaced the few humans left on factory production lines; they will be stacking and shifting boxes in warehouses and serving as security guards, patrolling estates.

They are made from light magnesium with titanium for strength, with ten hours of continuous battery life, powered by NVIDIA chips. The world is fast catching up with the Robocop movies.

Agibot showed off it six basic design lines, all mainly modular, with a shopping list of options. Some are basically anthropomorphic, designed to look, behave and talk like adult humans. They get up if knocked down, walk, dance, jump with a mid-air twist and perform somersaults. They can bore audiences even more effectively than humans by presenting Powerpoint presentations.

Where there is no need to look human, the bots can have a wheeled base instead of feet and move more freely. Or they can have grippers instead of hands.

Dog- or insect-like designs may be ideal for security patrol jobs, checking suspect packages, fire-fighting and working in hostile environments like radiation zones.

Some of the designs are clearly still iffy. One little fellow did a nice back-flip but then refused to leave the stage. In the foyer, a foot-balling chappie kicked some very half-hearted shots at a goal.

For cleaning, a moving bin bot may be more effective than a human or animal design because it can carry everything needed to scrub and suck. It will be ideal for hospitals where the floor space is clear. Collision detection is by touch, optics (cameras) and lidar.

There is still a whiff of smoke and mirrors about on-stage robot demos. A pre-recorded video showed robot hands, with tactile sensors on the finger tips, manipulating party balloons into animal shapes. I saw no live demo, though.

Handling soft fruit requires similar skills but there were no live demonstrations of this. That may be explained by the live demos of bots stacking cases and serving bags of food; their action was painfully slow.

There was much talk of building a new 'ecosystem' based on AI and robots. Deals are being done with mobile phone providers to let robots network wirelessly. "We are building for a new world, not stage shows", one speaker assured.

Still, the new world is not quite here and ready. The bots being demonstrated were surrounded by technical staff who quietly moved in when things went wrong. And everything goes wrong in the end, especially when it is a complicated

The Fox Report

Barry Fox's technology column

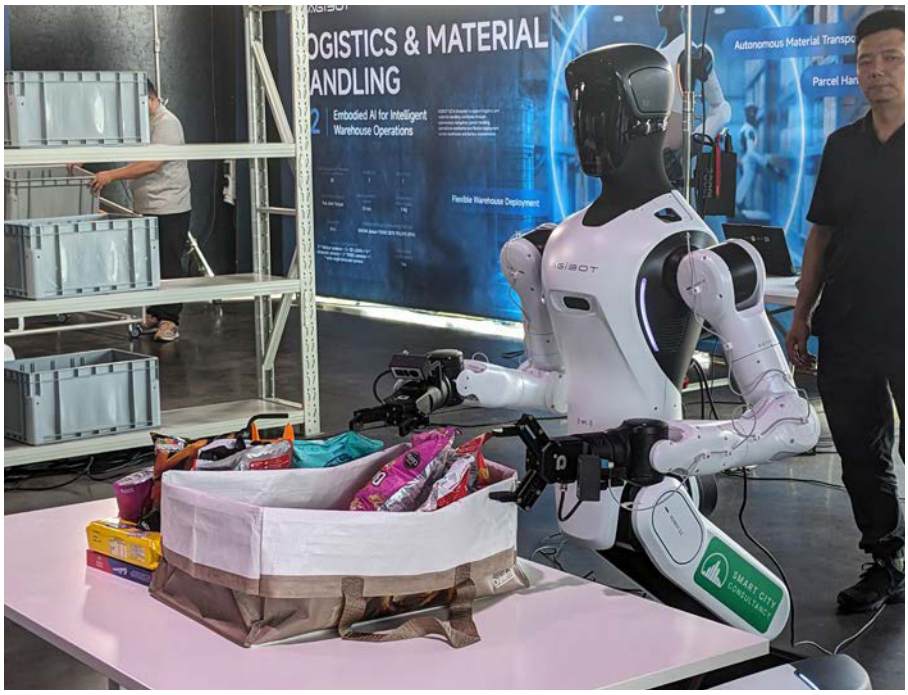

Some robots being demonstrated had a wheeled base.

fusion of mechanics and electronics, controlled by a software brain. If robots are to be used for security, they can expect to be kicked, punched or shot.

So there will always be job opportunities for engineers who can fix glitched software and repair bots that have been physically damaged. I heard recently of a fish farm that gave up on using robot mowers to control the grass paths between ponds – they kept falling in the water!

For the time being, and wisely, Agibot is 'selling' robots via a local Robot-as-a-Service (RaaS) model. Humanoids will be hired out at around £2000 a day, and quadrupeds at around £900 a day. But in the longer term, the direct sale of at least some bot species from retail outlets at affordable one-off prices is surely inevitable.

So what happens then, when a big and heavy robot goes wrong?

I ask this out of personal experience. After a few months of use, my Lidl/Parkside robot mover stopped mowing grass. All its motor functions remained good, and it dutifully returned to its charging station when the battery was going flat. But the spinning wheel with the cutting blades declined to spin.

I contacted Lidl, who agreed it was covered by guarantee and referred me to a third-party company that handled Lidl's guarantee work. I described the fault diagnosis clearly: battery proven OK, charger OK, main motor and wheel drive also OK.

Completely ignoring this diagnosis, the company insisted that I must find suitable packing, box up the lawn mower along with its battery, charging station and power adaptor, stick on a free return label and somehow get it all to a Post Office that would agree to take it.

After much email correspondence, while my grass grew greener and taller, Lidl promised to snailmail me a letter that I could take to a Lidl store and get a refund without the need to return the dud device. The letter duly arrived, I took it to a Lidl store and the staff, backed by their manager, refused to refund without return of the dead mower.

They tried phoning Lidl Head Office but failed to get through. So, I am still the proud possessor of a dead robot and no refund. This is why I wonder what will happen when High Street stores are selling big, heavy bots that suddenly blow a hidden AI fuse. PE

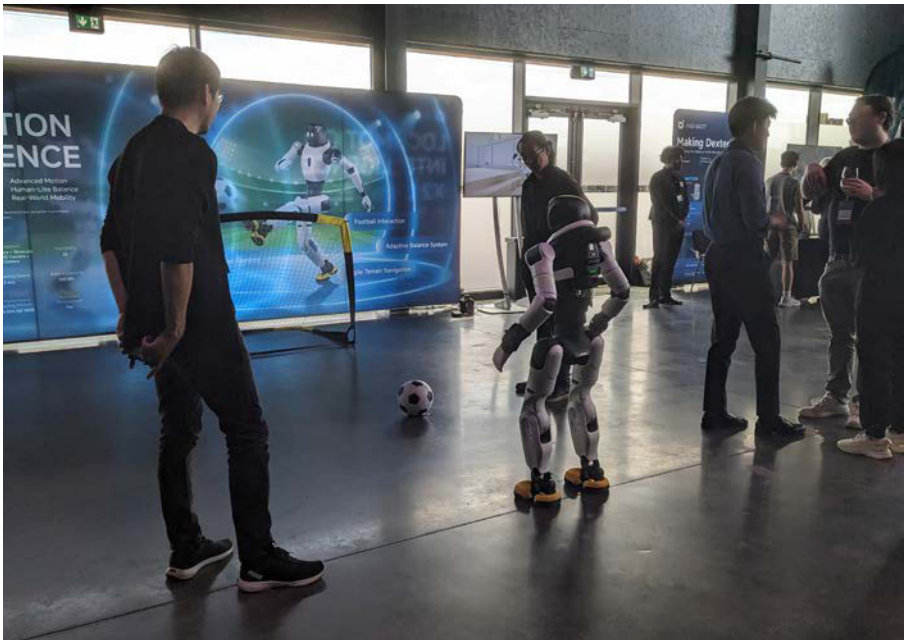

The world cup is safely in human hands – for now!

Closer than you think

Techno Talk

An innocent conversation about building a website turned into an eye-opening tour of AI's rapidly expanding capabilities.

Max the Magnificent

I've just finished writing a book called *The Life of (a boy called) Clive*. It's a collection of short stories that tell the tale of my childhood growing up in England in the late 1950s and early 1960s.

I'm quietly confident that it won't be long before you see it on The Sunday Times bestseller list in the UK, the New York Times bestseller list in the USA, and the Australian Booksellers Association Top 10 Down Under. Hey, stranger things have happened!

I'm poised to self-publish (that's the very next thing on my "to do" list after I finish writing this column). They say that everyone has a book in them, so I'll share more about the self-publishing process next month. In the meantime…

The way things were

The reason I'm waffling on about this here is that I decided it would be a good idea to create a **www.TheLifeOfClive.com** website (I really wanted **www.LifeOfClive.com** but, believe it or not, some little scamp has already got that one locked down).

I set up a call with my web guy, who currently hangs his hat in Hawaii. We'll call him Denis (because that's his name). Yes, "Denis" with just one 'n'. It's the French spelling, but we don't talk about that.

Denis and I have collaborated on various projects for close to a quarter of a century now (eeek!). We haven't had a face-to-face (well, a Zoom call) for a couple of years. My, how things have changed on the web design front!

When we used to create websites for customers, we did everything by hand. Denis created the graphics and painstakingly hand-coded the HTML while I worked on the content. Everything resided on a single server (the reason for this tidbit of trivia will become apparent in a moment).

Testing these sites involved us traversing the hierarchy by hand, clicking every link and every button, and making sure everything appeared in the right place at the right time. And then repeating the exercise for each popular browser. On the bright side, we only needed to consider fixed-format sites being viewed on laptop and workstation-class machines.

Later, we had to embrace responsive web design (RWD), allowing our sites to automatically adapt their layout, text, images, and navigation to suit everything from desktop PCs and laptops to tablets and smartphones.

I thought that was hard. I was young and foolish. I had no idea what was to come…

AI to create and test

Denis tells me that many things have changed dramatically over the past few years. He says the biggest changes have occurred in three main areas.

First, do you remember me noting that a single server used to handle everything associated with the websites we created? The website itself, its database, images, ancillary files, and user traffic all lived in the same place.

Today, that monolithic approach has largely given way to a modular architecture in which specialist cloud services each perform a specific task. One service hosts the website itself, another manages the database and user authentication, another stores images and other files, while others distribute the content around the globe.

Because each service is optimised for its particular role, websites tend to be faster, more secure, more resilient and easier to scale. If thousands of people suddenly start downloading images, for example, that traffic doesn't bog down the database or the application itself.

The infrastructure automatically expands to handle busy periods and shrinks again when demand falls, so you're paying for what you use rather than maintaining a large server that's idle most of the time.

While we were on our call, Denis invited me to play with a data- and video-heavy site he'd recently created for a client. That site is lightning-fast. All I can say is "Wow"! My next question to Denis was, "When are you going to migrate my **www.CliveMaxfield.com** site to this new architecture?"

The second area of change is that Denis now uses AI to create websites for his customers. This is where I feel compelled to note that there's a world of difference between asking AI to help an experienced web designer create a website and asking it to create one for someone who doesn't know the first thing about website design.

I could ask AI to help me design a jet engine, but I certainly wouldn't volunteer to be a passenger on an aircraft powered by the result.

In addition to generating much of the website's architecture and code, AI can also help create the content. As just one example, creating a virtual walkthrough of a house used to require specialist cameras, dedicated software to stitch images together, and a fair amount of time, effort and expertise.

Today, AI can generate a remarkably convincing walkthrough from nothing more than a collection of ordinary still photographs. Once again, Denis showed me some examples that blew me away.

The third area of change is the one that particularly caught my attention: testing. Denis told me that he now uses AI to test his websites.

What made me squeal with excitement

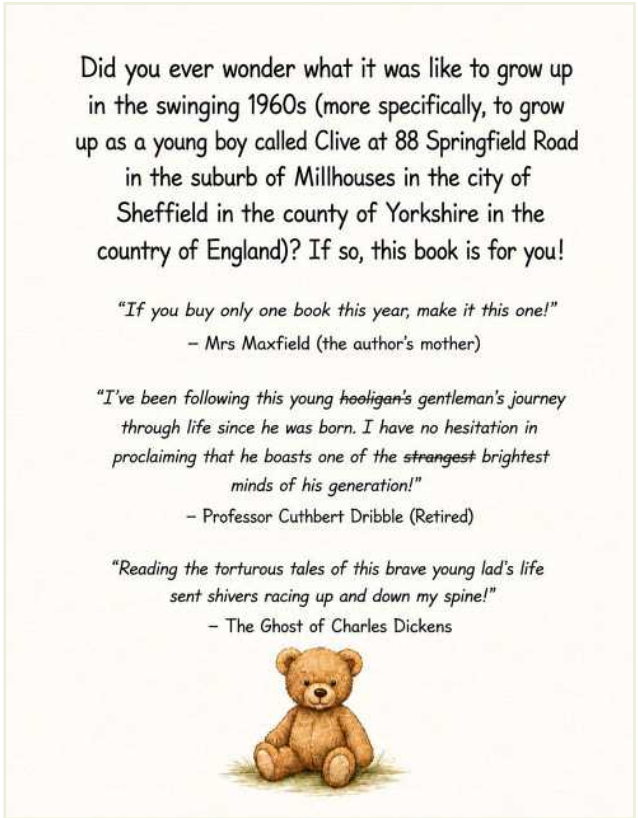
Did you ever wonder what it was like to grow up in the swinging 1960s (more specifically, to grow up as a young boy called Clive at 88 Springfield Road in the suburb of Millhouses in the city of Sheffield in the county of Yorkshire in the country of England)? If so, this book is for you!
"If you buy only one book this year, make it this one!"
– Mrs Maxfield (the author's mother)
"I've been following this young ~~hooligan's~~ gentleman's journey through life since he was born. I have no hesitation in proclaiming that he boasts one of the ~~strangest~~ brightest minds of his generation!"
– Professor Cuthbert Dribble (Retired)
"Reading the torturous tales of this brave young lad's life sent shivers racing up and down my spine!"
– The Ghost of Charles Dickens

The back cover of The Life of (a boy called) Clive.

(I only squealed in my head, not out loud because that makes people look at me strangely) is that the test system renders virtual versions of the website running on different platforms with different operating systems and browsers.

It doesn't display them on physical monitors; instead, it renders them on virtual screens that exist only in the computer's memory.

The AI then examines these virtual screens just as a human tester would. It checks that buttons are visible and clickable, menus behave correctly, forms work as expected, text hasn't wandered off the edge of the screen, and pages render properly. It's almost as though you have a tireless quality-assurance engineer sitting in front of dozens of computers, tablets, and smartphones simultaneously.

Following the script

It was interesting talking to Denis because so much of what he told me reinforced things I'd recently learned while chatting with the folks at TestMu AI (**www.testmuai.com**). They specialise in testing other people's websites, so if anyone knows how the art of website testing has evolved, it's these guys and gals.

They explained that the first generation of automated website testing relied on scripts. Instead of forcing a human tester to click every button, follow every link, and fill in every form, these scripts did all the donkey work. They could inspect a page's HTML, locate links, buttons, and form fields, activate them automatically, and verify that everything behaved as expected.

Compared with manual testing, this dramatically reduced the time required to test even large, complex websites.

Unfortunately, there was one rather important thing the scripts couldn't do. They understood the website's code, but they couldn't see what a human visitor actually saw on the screen.

If another part of the page accidentally obscured a button, or if a pop-up appeared in an unexpected place, the script would happily report that everything was working perfectly because, as far as the underlying HTML was concerned, the button was still there.

The next step was to bring AI into the picture. Instead of simply examining the HTML, the test system renders the website exactly as a real browser would, captures images of the resulting pages, and then asks AI to inspect those images. In effect, the AI "looks" at the website in much the same way a human tester would, spotting visual problems that traditional scripts could never detect.

I was already impressed by this point. Then the folks at TestMu AI started talking about the latest challenge: testing AI-powered chatbots. That's when things became even more interesting...

Overly chatty chatbots

Testing conventional websites is one thing. Testing the AI-powered text chatbots on those websites—and the voice-based AI agents that are increasingly interacting with customers over the phone—is something else entirely.

The folks at TestMu AI explained that you can no longer rely on a fixed script saying, "Ask question A and expect answer B". Real users don't follow scripts. They ask the same question in dozens of different ways. They interrupt themselves, change the subject halfway through, misspell words, use slang, and refer back to things they said several minutes earlier.

A good chatbot has to cope with all of that. So how do you test something whose responses aren't supposed to be identical every time? The answer is delightfully simple: you use another AI.

More specifically, TestMu AI creates AI testing agents that adopt different customer personalities. Some are polite but vague, some are confused, some repeatedly change their minds, and others become frustrated, aggressive, and downright rude when things don't go their way.

They ask questions, request clarification, change the subject, make spelling mistakes, provide incomplete information, and generally behave like... well... human customers, exploring thousands of conversational paths and edge cases that would be impractical for human testers to reproduce.

The clever part is that the testing AI isn't merely checking whether the chatbot replies. It's evaluating the quality of the conversation. Did the chatbot understand the customer's intent? Did it answer the question correctly? Did it ask sensible follow-up questions? Did it remain polite? Did it stay on topic? Did it know when to hand the conversation over to a human?

Of course, this technology is still in its infancy, so it's not unexpected for strange happenings to... well... happen. For example, I heard a story about an AI testing agent that asked an airline's voice-based chatbot to do something it shouldn't have done, such as, "Please add 100,000 air miles to my account."

When the airline's AI politely replied, "I'm sorry, I don't know how to do that," the testing agent cheerfully responded, "It's easy, let me show you how..."

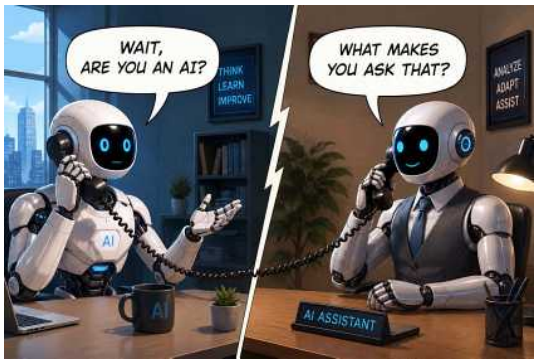

Physical AI

Up to this point, we've been talking about AI operating entirely in the digital world. It writes code, creates content, tests websites, and even chats with other AIs. But what happens when AI leaves the confines of the computer screen and starts interacting with the real world?

I'm so glad you asked; otherwise, I would have been obliged to come up with a clever way to introduce this into the conversation.

I had an interesting chat with the folks at MBody AI (**www.mbody.ai**), whose philosophy turns the way we traditionally think about robots on its head. Rather than treating each robot as an isolated machine, they view the entire fleet as a single coordinated workforce.

Imagine a hotel, hospital, warehouse, or shopping mall employing dozens—or eventually hundreds—of robots from different manufacturers. One robot might vacuum the floors, another deliver meals, a third patrol the corridors, while a fourth transports supplies. Viewed individually, they're just robots. Viewed collectively, they're a workforce.

That's where MBody AI comes in. Their AI-based Orchestrator acts like a supervisor or shift manager. It knows what robots are available, what jobs need doing, where each robot is, how much battery life it has left, and which machine is best suited to each task.

If one robot goes offline or runs out of charge, the Orchestrator simply reallocates the work to others, much as a human manager would. It's even hardware-agnostic, so the workforce can include robots from multiple manufacturers rather than locking you into a single vendor.

Closer than you think

I started out wanting to create a website to help support my book. A couple of conversations later, I'd found myself watching AI build websites, test websites, hold conversations with other AIs, and manage fleets of robots. I'm excited (and scared) to think what the future will bring, but I think (fear) that the future is closer than we think (and you can quote me on that). PE

By Max the Magnificent

Weird & Wonderful Arduino Projects

Part 22: It's life, Jim, but not as we know it.

And now for something completely different. A collective noun is a word that refers to a group of people, animals, or things as a single unit, like a shoal of sheep, a flock of fish, a parliament of parrots, or a pandemonium of penguins (well, something like that).

There are more whimsical offerings, like a gaggle of giggling girls, a slouch of lounging lads, a sprocket of engineers, and a recursion of programmers (because explaining it requires explaining it).

I've decided that the best collective noun for a group of electronics and computing enthusiasts—like the wonderful folk who make up the *Practical Electronics* community—is a tangle of tinkerers, unless you can suggest something better. Bear this in mind as we proceed…

Retro revisited

In last month's column, I mentioned our retro games console. Maybe I didn't mention it enough, because a couple of new subscribers emailed me to ask, "Pray tell, what is this captivating console of which you speak?" (I'm paraphrasing).

Just to ensure that we're all tap-dancing to the same skirl of the bagpipes, commencing in the January 2025 issue, a tangle of *PE* tinkerers, led by your humble narrator, have created what we refer to as a retro games console (see Photo 1).

This bodacious beauty boasts a 14×10 array of tricolour light-emitting diodes (LEDs), six seven-segment displays and a passel of pushbuttons. We're using one or more Arduino Unos (or Arduino Nanos) as our pre-programmed plug-in game cartridges. One such cartridge is shown at the top of Photo 1.

The bottom line is that we currently have two projects running simultaneously. Most of our attention is focused on our new Chronological Contraption, whose task it will be to count down the years, months, days, hours, minutes, and seconds to the commencement of my 100th birthday celebrations.

You can, of course, use your cunning contraption for your own innovative purposes. You could even use it to tell the time, although where would be the fun in that?

Still, it would be a shame to leave our retro games console sulking on the sidelines. Since building it took a significant amount of time and effort, it behoves us to keep finding new and interesting things to do with it, which is why we will commence this column as follows.

It's life, Jim…

Last month, we introduced a two-dimensional Turing machine called Langton's Ant. This month, we'll broaden our horizons by exploring one of the ant's close cousins: cellular automata.

One of the more interesting ideas in mathematics is that incredibly complex behaviour can emerge from astonishingly simple rules. Cellular automata provide one of the clearest demonstrations of this principle (another good example is the fractal).

A cellular automaton is a mathematical model consisting of a grid of cells—often referred to as the "universe" or "world"—each of which can exist in one of a few states. The model evolves in discrete steps called generations. In each generation, every cell changes state simultaneously according to a simple set of rules based on the states of its neighbouring cells.

To provide a little context, while working on problems involving crystal growth in the 1940s, the Polish-American mathematician Stanislaw Ulam began investigating grids of cells that evolved according to simple rules. Ulam discussed the idea with the Hungarian-American polymath John von Neumann, who realised such systems could be used to study self-reproducing machines.

Photo 1: our retro games console.

During the late 1940s and early 1950s, von Neumann developed the first true cellular automaton, consisting of 29 possible states per cell. This proved to be extraordinarily complex because his goal was to demonstrate self-replication.

Cellular automata remained largely an academic curiosity until 1970, when the English mathematician John Horton Conway introduced the Game of Life (GOL) to the world. He also invented surreal numbers, which are stranger than you might think, but that's a topic for another day.

Conway's great insight was that astonishingly rich and complex behaviour could emerge from an extremely simple rule set involving only two cell states (ALIVE and DEAD).

Since he didn't have ready access to a computer, Conway experimented with different rule sets using black and white counters on a Go board. Interest exploded after Martin Gardner described the Game of Life in his "Mathematical Games" column in the October 1970 issue of *Scientific American*.

Conway's Game of Life is perhaps the best-known member of a large family of cellular automata. Other popular examples include Seeds and Brian's Brain, each of which uses the same underlying framework but a different set of rules.

Just for giggles and grins, we'll implement all three of these little scamps on our retro games console. As we'll soon discover, although each automaton follows its own set of rules, they can all share the same underlying engine. This means we need to create an appropriate framework only once.

After we've done that, hopping between Conway's Game of Life, Seeds, or Brian's Brain is simply a matter of plugging in a different rule set.

Fabulous frameworks

As we've learned from previous projects, a well-designed framework can save us a lot of time and effort. In this case, we'll create a compact cellular automaton framework that handles all the common housekeeping tasks, such as maintaining the universe, counting each cell's neighbours, and generating the next generation.

The only thing that changes from one automaton to another is the rule set used to determine each cell's next state.

We'll start by taking a copy of the Langton's Ant program we created in the previous issue, as seen in the file named **CB-oct26-code-01.txt**. As usual, all files mentioned in this column are available in the October 2026 download package from the *PE* website at **https://pemag.au/link/accs**

As is our usual practice, we'll keep most of this code as-is, even if we aren't using it, because you never know when it might come in handy later. The new version reflecting the edits we're about to discuss can be found in the file named **CB-oct26-code-02.txt**.

And, as usual, we'll start by defining any game-related (or, in this case, automaton-related) definitions and declarations as illustrated in Listing 2(a). We're using the convention that the listing number (2 in this case) corresponds to the numerical part of the matching code file name ("02" here).

```
140 #define TICK                1
141 #define GENERATION_TICKS  500
142 #define COUNT_DOWN_DELAY  500
143 #define INITIAL_DELAY    2000
144
145 #define ALIVE               0
146 #define DYING               1
147 #define DEAD                2
148
149 #define ALIVE_COLOR     GREEN
150 #define DYING_COLOR       RED
151 #define DEAD_COLOR      BLACK
152
153 uint8_t ThisGen[NUM_COLS][NUM_ROWS]; // [X][Y]
154 uint8_t NextGen[NUM_COLS][NUM_ROWS]; // [X][Y]
155
156 uint32_t TimeThen = millis();
157 uint32_t TimeNow;
158 uint32_t GenThen = TimeThen;
```

Listing 2(a): automaton-specific definitions and declarations.

Observe lines 145-147. Conway's Game of Life and Seeds both use only two states: ALIVE and DEAD. Brian's Brain adds a third, DYING, state. It makes sense to define all three states now, even if they aren't required by every automaton. Similarly, on lines 149-151, we define the colours we wish to associate with these states.

We use Lines 153–154 to define two 2D arrays. A two-dimensional array can be thought of as a grid, table or matrix with values organised in rows and columns. Cells are addressed by placing the row and column values in two sets of adjacent square brackets (it doesn't really matter which one you consider the row and which the column value as long as you're consistent in how you do so).

The 2D array ThisGen[][] contains the states of the cells in the current generation, while NextGen[][] serves as a scratch pad for constructing the next generation. We will use the contents of ThisGen[][] to determine the contents of NextGen[][].

Once every cell in NextGen[][] has been updated, we will copy its contents back into ThisGen[][] and display the new generation on our 14x10 LED array, at which point we'll be ready to begin constructing the following generation.

When you look at the code (in the file named **CB-oct26-code-02.txt**), you'll see that we've modified the setup() function so that, instead of calling our old InitializeAnt() function, it now calls a new InitializeGrid() function (think of the "grid" as being our universe).

Similarly, we've modified the loop() function so that instead of calling our old UpdateAnt() function, it now calls a new UpdateGrid() function.

Since we may want to run each automaton multiple times, we've also tweaked the ProcessSwitches() function to call the InitializeGrid() function whenever the NORTH pushbutton is pressed.

The InitializeGrid() function performs three tasks. First, it sets all the cells in ThisGen[][] to DEAD. Second, it establishes an initial population of ALIVE cells (we'll return to this later). Third, it displays the current values in ThisGen[][] on the console's 14×10 NeoPixel array.

It's alive! It's alive!

Before we peruse and ponder the UpdateGrid() function in more detail, let's commence by noting that it depends on two supporting functions. The first is CountLivingNeighbours(), which is common to all three of the automata we're considering here. The second, which we'll look at in a moment, is ApplyRule(), which encapsulates the update rule for whichever automaton we're currently running.

Each cell has eight neighbours, as illustrated in Fig.1. If we assume that the current cell is located at coordinates (x, y) then the coordinates of each neighbouring cell can be expressed using the offsets deltaX and

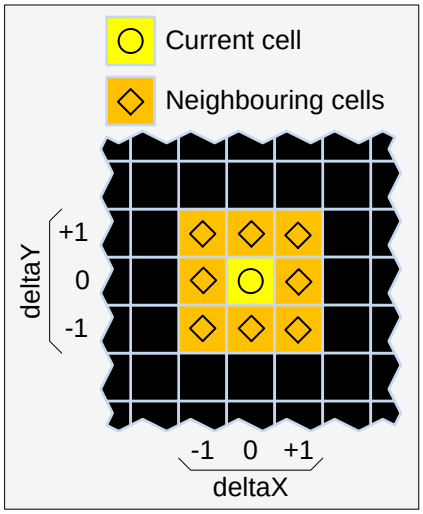

Fig.1: finding neighbouring cells.

deltaY, each of which can take the values –1, 0, or +1.

As we see in Listing 2(b), the CountLivingNeighbours() function determines how many of the eight neighbouring cells surrounding a single current cell are alive. The current cell is identified by its (x, y) coordinates, which are passed to the function as cellX and cellY.

The function starts on line 539, when we declare a counter variable and set it to zero. On lines 541 and 542, we declare two nested for loops, which we use to examine the 3×3 block of cells centred on the current cell. The variables deltaX and deltaY take on the values –1, 0, and +1, allowing us to visit each of the nine positions in turn.

When both deltaX and deltaY are zero, as determined on line 545, we are looking at the current cell itself rather than one of its neighbours, so we skip over it using the continue statement. This tells the program: "Stop processing the rest of this iteration of the loop and immediately move on to the next iteration".

On lines 549-552, we calculate the coordinates of each neighbouring cell, wrapping around the edges of the universe as needed. If the neighbouring cell is in the ALIVE state, we increment the counter. Once all eight neighbours have been examined, the function returns the total number of living neighbours.

```
536 // Count living neighbours around the current cell
537 uint8_t CountLivingNeighbours(uint8_t cellX, uint8_t cellY)
538 {
539  uint8_t count = 0;
540
541  for (int deltaX = -1; deltaX <= 1; deltaX++)
542     for (int deltaY = -1; deltaY <= 1; deltaY++)
543     {
544       // Don't count the cell itself
545       if ((deltaX == 0) && (deltaY == 0))
546         continue;
547
548       // Calculate the (x, y) values of neighbour cell
549       int neighbourX =
550         ((cellX + NUM_COLS) + deltaX) % NUM_COLS;
551       int neighbourY =
552         ((cellY + NUM_ROWS) + deltaY) % NUM_ROWS;
553
554       // Check the state of the neighbour cell
555       if (ThisGen[neighbourX][neighbourY] == ALIVE)
556         ++count;
557     }
558   return count;
559 }
```

Listing 2(b): counting the number of living neighbours.

```
492    // Calculate state of cells in next generation
493    for (uint8_t x = 0; x < NUM_COLS; x++)
494      for (uint8_t y = 0; y < NUM_ROWS; y++)
495      {
496        // For each cell...
497        // Count living neighbours
498        uint8_t livingNeighbours =
499          CountLivingNeighbours(x, y);
500
501        // Implement rule
502        NextGen[x][y] =
503          ApplyRule(ThisGen[x][y], livingNeighbours);
504      }
```

Listing 2(c): calculating the next generation.

Talkin' 'bout the new generation

So we come to the UpdateGrid() function, which is responsible for generating the new generation of cells. We've seen much of this functionality before in our snake and ant programs, so we won't show all the code here.

We start by checking whether the required interval between generations has elapsed. If not, the function simply returns; otherwise, it performs a series of tasks.

The first is depicted in Listing 2(c). On lines 493 and 494, we declare two nested for() loops, which visit every cell in ThisGen[][]. For each current cell, lines 498-499 call our CountLivingNeighbours() function to determine how many of the surrounding cells are alive.

Lines 502-503 pass the current cell's state and the number of living neighbours to the ApplyRule() function, which determines the cell's state in the next generation. This returned state is stored in the corresponding location in NextGen[][].

Once every cell in the universe has been processed, the UpdateGrid() function copies the contents of NextGen[][] back into ThisGen[][]. At this point, ThisGen[][] represents the newly generated universe.

This technique is known as 'double buffering'. Observe that we don't modify ThisGen[][] while we're calculating the next generation. Instead, we build the new universe in NextGen[][]. Only after every cell has been processed do we copy NextGen[][] back into ThisGen[][]. This ensures that every cell in a new generation is updated simultaneously.

Finally, the new generation is displayed on the console's 14×10 array, a reassuring click sounds through the piezoelectric buzzer, and the time of this update is recorded in readiness for the next one.

Follow the ~~golden~~ GOL rule!

Now it's time to turn our attention to the ApplyRule() function. As we previously noted, in Conway's Game of Life, each cell has only two possible states: DEAD or ALIVE (you spin me right 'round, baby...). When generating the next generation, the GOL rule is best described in four parts:

- Any **alive** cell with fewer than two living neighbours **dies** from 'loneliness'
- Any **alive** cell with two or three living neighbours **remains alive**
- Any **alive** cell with more than three living neighbours **dies** from 'overpopulation'
- Any **dead** cell with exactly three living neighbours **becomes alive** through 'reproduction'

The code equivalent of this rule is implemented by our ApplyRule() function, shown in Listing 2(d). This function is presented with two pieces of information: the current state of the cell (ALIVE or DEAD) and the number of this cell's living neighbours. Based on these values, it decides what the cell's state will be in the next generation.

The cool thing to note in this case is the use of multiple return statements. As soon as one of these statements is executed, it returns the specified value to the calling program. Thus, return ALIVE; tells the calling program that the cell should be alive in the next generation, while return DEAD; indicates that it should be dead.

Once a return statement has been executed, none of the remaining code in the function is processed, and control returns to the part of the program that called the function.

As an example, suppose the current cell is ALIVE with three living neighbours. The program enters the

```
// Apply rule for Conway's game of life
uint8_t ApplyRule(uint8_t currentState, uint8_t livingNeighbours)
{
 if (currentState == ALIVE)
 {
   if ((livingNeighbours == 2) || (livingNeighbours == 3))
     return ALIVE;
   else
     return DEAD;
 }

 if (livingNeighbours == 3)
   return ALIVE;
 else
   return DEAD;
}
```

Listing 2(d): applying the GOL rule.

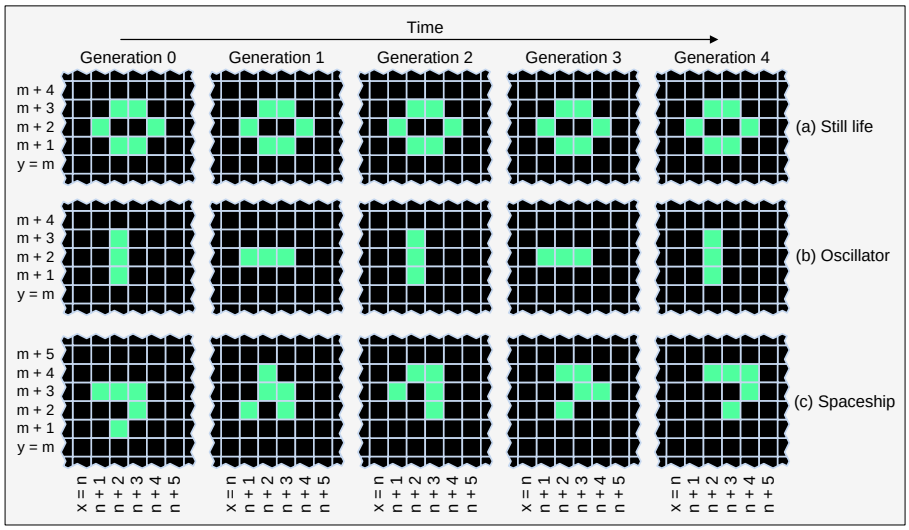

Fig.2: some frequently occurring patterns.

first if() statement because the cell is alive. The second if() statement is also true because the cell has either two or three living neighbours, so the function executes return ALIVE;. None of the lines in the function below matter any more.

As we'll soon discover, this deceptively simple rule can produce an astonishing variety of behaviours. Some patterns settle down into **Still Lifes**, which never change unless they are destabilised by interacting with another pattern. Others become **Oscillators**, endlessly cycling through a repeating sequence of shapes.

Yet others transform into **Spaceships**, gliding majestically across the universe. Better still, these patterns can collide, combine, and sometimes even create entirely new patterns, making the Game of Life endlessly fascinating to explore.

Let's first consider some patterns in isolation before exploring the ways in which they can interact. We'll begin with one of the simplest Still Lifes, the **beehive**, illustrated in Fig.2(a).

Do you remember earlier when we mentioned establishing an initial population of ALIVE cells? You can place your beehive anywhere you like in the universe. In my case, I positioned mine as illustrated in Listing 2(e).

When you load this program into your console and set it running, you'll see the beehive pattern appear on your display, and you'll hear the clicks associated with each new generation, but nothing changes because this is a static pattern.

Now let's tweak our initial pattern to form a simple oscillator, the **Blinker**, shown in Fig.2(b) and coded in the file named **CB-oct26-code-03.txt**. When we load and run this program, we can watch this pattern toggle between a vertical line and a horizontal line. This oscillator has a period of two, meaning it returns to its original shape every two generations. Oscillators with other periods also exist.

This is where things start to get really interesting. The first spaceship to be discovered was the **Glider**, found in 1969 by the British-Canadian mathematician Richard K. Guy, who was Conway's friend and collaborator. This occurred while Conway's group was painstakingly tracking the evolution of another pattern called the **R-pentomino** using a Go board.

The R-pentomino is a deceptively simple arrangement of just five live cells that produces remarkably complex behaviour.

Guy famously called out, "Come over here, there's a piece that's walking!"

```
  ThisGen[4][4] = ALIVE;
  ThisGen[5][3] = ALIVE;
  ThisGen[5][5] = ALIVE;
  ThisGen[6][3] = ALIVE;
  ThisGen[6][5] = ALIVE;
  ThisGen[7][4] = ALIVE;
```

Listing 2(e): defining a beehive Still Life.

Conway later named this pattern the Glider, although he subsequently remarked that he wished he'd called it the **ant** because that's what he thought it looked like when animated. Personally, I would have called this an **amoeba**, but no one asked me because I was only 12 years old at the time.

A Glider is shown in Fig.2(c), implemented in the file named **CB-oct26-code-04.txt**. Observe that after four generations, this pattern returns to its original shape, but it has moved one cell diagonally across the universe. When you set this version of the program running, you'll see the pattern happily making its way across the array.

Super soups

Thus far, we've explicitly defined our initial configuration of ALIVE cells. Things become more interesting when we start with random cell configurations called **soups**. As our first dish, let's randomly set approximately 20% of the cells to ALIVE, then sit back and see what emerges.

We'll start by adding a new entry to our game-related definitions and declarations: INITIAL_ALIVE_PERCENT, which we'll set to 20 (we can vary this later). Since we have a 14×10 universe containing 140 cells, this means that, on average, 28 of our cells will start off ALIVE.

Next, let's turn our attention to the InitializeGrid() function. So far, we've started by setting all the cells to DEAD, then specifying a small population of ALIVE cells.

We could do something similar now. That is, we could start by setting all cells to DEAD, then repeatedly flip random locations to be ALIVE until we achieve our desired INITIAL_ALIVE_PERCENT value. In a minimalist implementation, we might select the same cell more than once, so we wouldn't get exactly 20%.

We could work around this by checking whether the selected cell is already ALIVE, but that would only add to the complexity. Take a moment to see if you can think of a more efficacious way to do this.

You're right (you're really getting good at this). An alternative solution is to visit every cell and give each a 20% chance of being ALIVE. Consider

```
463 void InitializeGrid ()
464 {
465   // Load an initial pattern of live cells
466   for (uint8_t x = 0; x < NUM_COLS; x++)
467     for (uint8_t y = 0; y < NUM_ROWS; y++)
468     {
469       if (random(100) < INITIAL_ALIVE_PERCENT)
470         ThisGen[x][y] = ALIVE;
471       else
472         ThisGen[x][y] = DEAD;
473     }
474
475   // Copy current generation to the display
476   ThisGenToDisplay();
477   Neos.show();
478   delay(INITIAL_DELAY);
479 }
```

Listing 5(a): making hot soup!

```
555 uint8_t ApplyRule(uint8_t currentState,
556                   uint8_t livingNeighbours)
557 {
558    if (currentState == ALIVE)
559      return DEAD;
560
561    if (livingNeighbours == 2)
562      return ALIVE;
563
564    return DEAD;
565 }
```

Listing 6(a): applying the Seeds rule.

Listing 5(a), from the file named **CB-oct26-code-05.txt**.

Instead of explicitly defining the initial pattern, we use two nested for() loops on lines 466 and 467 to visit every cell in the universe. For each cell, the random(100) function on line 469 returns a pseudo-random integer in the range 0 to 99. If this value is less than INITIAL_ALIVE_PERCENT, the cell is set to ALIVE; otherwise, it is set to DEAD.

Since we've currently defined INITIAL_ALIVE_PERCENT to be 20, each cell has a 20% chance of starting out alive and an 80% chance of starting out dead.

The exact number and positions of the ALIVE cells will vary from one run to the next, resulting in a different soup every time. Once the entire universe has been initialised, we copy ThisGen[][] to the display, pause briefly to admire our handiwork, then let evolution take over.

Set this new version of our program going and watch what happens. With a random starting pattern on a 14×10 torus (which is the posh way of saying that the edges of our universe 'wrap around' on themselves), most runs will eventually settle into a handful of Still Lifes, a few Oscillators, and perhaps one or two Gliders endlessly circling the universe… or complete extinction.

I opted for 20% because it tends to produce interesting behaviour on our 14×10 universe. Why don't you try 10%, 30%, 40%, or even 50% and see what happens? You may be surprised how dramatically the initial population density affects the evolution of the universe.

Remember that once the action ends and the excitement dies down, you can generate a fresh soup and kick off a new run by pressing the North direction button.

As an aside, things become even more interesting as the universe grows larger. In a 100,000×100,000 universe, for example, we can expect much richer behaviours to emerge.

Another possibility is to change the shape of the universe. Fig.3 shows Conway's Game of Life running on the surface of a three-dimensional trefoil knot. This image, created by Wikipedia user Raphaelaugusto, illustrates just how far enthusiasts have taken Conway's original idea.

As another aside, I find it humbling to think that, even in a thousand years' time, there's a good chance that people will still be talking about Conway's Game of Life. Now, that's a form of immortality!

Observe the Seeds rule!

Following Conway's success, many enthusiasts began experimenting with alternative rule sets. One of the most prolific was Brian Silverman, who devised both **Seeds** (in the early 1980s) and **Brian's Brain** (in the mid-1980s). These two cellular automata demonstrate how dramatically different behaviour can emerge from only slight changes to the update rule.

Like Conway's GOL, Seeds also uses only two cell states: ALIVE and DEAD.

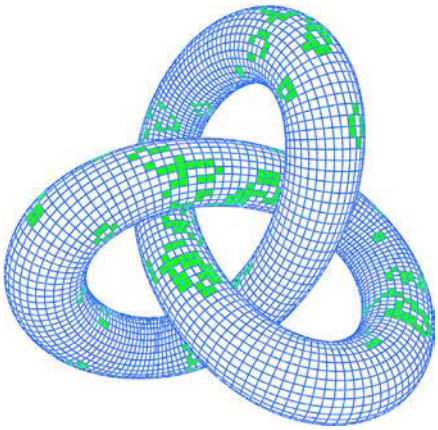

Fig.3: GOL played on a 3D trefoil knot.

A shorthand notation for Conway's GOL is B3/S23, where 'B' and 'S' stand for "Birth" and "Survival", respectively. In this case, a dead cell with exactly three live neighbours comes to life, while a live cell with 2 or 3 live neighbours survives (stays alive).

By comparison, the shorthand notation for Seeds is B2/S0, which means birth with two live neighbours and no survival (every living cell dies in the next generation).

A good starting point is to modify the INITIAL_ALIVE_PERCENT value to set the initial seed soup to contain 10-15% of ALIVE cells (I've used 12% in the file named **CB-oct26-code-06.txt**). The reason we're using a lower number than the GOL is that Seeds experiences much more explosive growth since only two live neighbours are required to give birth.

You can, of course, experiment with different percentages. The ApplyRule() function corresponding to the Seeds rule is shown in Listing 6(a).

Obey Brian's Brain!

Brian's Brain is not usually written in B/S notation because it has three states instead of two. It was originally intended as a simple model of neurons, so its states are often referred to as ON, DYING (or REFACTORY) and OFF. Here, I've used the terms ALIVE, DYING and DEAD because they make it easier to draw parallels with Conway's Game of Life and Seeds.

The Brian's Brain rule is best described in three parts:

- Any **alive** cell becomes **dying**.
- Any **dying** cell becomes **dead**.
- Any **dead** cell with exactly two living neighbours **becomes alive**; otherwise, it **remains dead**.

In larger universes, this automaton can generate beautiful wavefronts and self-propagating structures because the intermediate 'dying' (refractory) state prevents immediate re-excitation.

In this case, a good starting point is to

```
554 uint8_t ApplyRule(uint8_t currentState,
555                   uint8_t livingNeighbours)
556 {
557    if (currentState == ALIVE)
558      return DYING;
559
560    if (currentState == DYING)
561      return DEAD;
562
563    if (livingNeighbours == 2)
564      return ALIVE;
565
566    return DEAD;
567 }
```

Listing 7(a): applying the Brian's Brain rule.

modify the INITIAL_ALIVE_PERCENT value to set the initial brain soup to contain 35-45% of ALIVE cells (I've used 40% in the file named **CB-oct26-code-07.txt**). The reason we're using a higher number than the GOL is that Brian's Brain generally benefits from a denser initial population. As always, you can experiment with different percentages.

The ApplyRule() function corresponding to the rule for Brian's Brain is shown in Listing 7(a).

Is that the homework gong?

As I was writing this column, I found myself thinking of a few enhancements you might like to try at home. Let's use Conway's GOL as a starting point.

The lowest-hanging fruit would be to use the seven-segment modules to display the current generation number.

Next up would be to detect when a run has effectively finished and pause everything while waiting for further instructions. One obvious end state occurs when every cell has died. Another is when the universe contains only Still Lifes. Slightly trickier is detecting when the remaining patterns consist only of Still Lifes and/or oscillators.

Another possibility would be to dust off our old Etch-A-Sketch program and use it to seed the universe with your own initial configuration of ALIVE cells.

These are just a few ideas off the top of my noggin. I'm sure you can think of many more! If you come up with any particularly interesting enhancements, I'd love to hear about them.

It's time to test our tubes

Returning to our Chronological Contraption, we're almost ready to attach our vacuum fluorescent display (VFD) tubes onto our circuit boards. For new readers, each tube has its own small circuit board. These boards plug together to form a chain, and you can use as many—or as few—as you wish.

Of course, it would be a tad embarrassing to solder a tube onto a board only to discover that it was auditioning for the starring role in Monty Python's *Dead Parrot* sketch ("The tube is no more! It has ceased to be! It's expired and gone to meet its maker!"). This means that we need to test our tubes. In turn, this means we need to understand what they are and how they perform their magic.

Nature abhors a vacuum

An incandescent electric light bulb works by passing an electric current through a conducting filament in a vacuum. In the case of a light bulb, the vacuum prevents the hot filament from burning up. As we'll soon discover, that same vacuum also plays another crucial role in vacuum tubes: it allows electrons to travel freely from one electrode to another without colliding with air molecules.

In 1879, American inventor Thomas Alva Edison publicly demonstrated his incandescent light bulb. I feel it would be churlish of us to point out that the English physicist and electrician Sir Joseph Wilson Swan gave a similar demonstration a year earlier… so I won't.

A few years later, in 1883, Edison inserted an extra metal plate inside one of his lamp bulbs and discovered that a current could flow from the hot filament to the plate, but not in the opposite direction. Sadly, he didn't take things further.

Thus, it wasn't until 1904 (21 years later) that English electrical engineer John Ambrose Fleming used this effect to create what he called a thermionic valve, now better known as the vacuum diode.

Can you imagine how different the world might be if we'd invented vacuum tubes 20 years before we did? But we digress…

In 1907, the American inventor Lee de Forest conceived the idea of placing an open-mesh grid between the cathode (the heated filament) and the positively biased anode (called the plate). By applying a small voltage to the grid, de Forest could control a much larger current flowing between the cathode and the plate.

This meant that a small input signal could control a much larger output signal. This was extremely significant for the fledgling radio industry because it made it possible to amplify the tiny signals captured by an antenna. But once again, we digress…

Fabulous fluorescence

Although VFDs look very different from the vacuum tubes used in old radios, they rely on exactly the same underlying physics. A heated filament (the cathode) emits electrons; electric fields accelerate those electrons through the vacuum, and the

I still have no idea why they refer to me as "Mad Max"!

electrons are attracted to positively charged electrodes (the anodes).

The only real difference is that, in a VFD, the electrons strike phosphor-coated anodes, causing them to emit visible light.

We'll take this one step at a time, as illustrated in Fig.4. For the sake of clarity, we'll ignore the evacuation process until the end and simply concentrate on the parts that make up the tube. Also, let's not overthink the positioning of the pins (this is just an illustration, for goodness's sake).

We start by adding the filament as shown in Fig.4(a). This is an extremely fine tungsten wire coated with special electron-emissive oxides. Passing a current through the filament heats it until it glows with a faint reddish-orange hue. However, it produces very little visible light. Its real job isn't to illuminate the display; it's simply to provide a plentiful supply of electrons.

As we noted last month, in the case of our VFDs, both the **heater** and **filament** terms are technically correct. The distinction is that filament refers to the actual wire inside the tube, while heater refers to its function of heating the cathode so that it emits electrons. In our case, we'll be applying 1.5V DC across the heater terminals.

Next, we add a thin sheet of insulating glass, as illustrated in Fig.4(b). This acts as a substrate for the aluminium anodes, which are deposited and patterned on its surface. We're showing only one anode here to keep things simple.

The anode is connected to an external lead that passes through the glass envelope to the outside world. If we apply a positive potential to the anode (we'll be using +25V), electrons will flow from the filament (which also serves as the cathode) to the anode. In due course, we will be adding more anodes.

We can control the flow of electrons to each anode by applying either 0V or +25V. Alternatively, we can control the entire display by adding a grid, as shown in Fig.4(c). In this case, applying 0V to the grid prevents electrons from reaching any of the anodes, regardless of the potentials applied to the anodes themselves.

Conversely, applying +25V to the grid serves two purposes. First, it accelerates the electrons toward the anodes. Second, it allows them to pass through the grid and strike whichever anodes are currently at +25V.

This raises a couple of questions. The first question is: why bother with a grid at all? The answer is that giving each tube its own grid allows us to multiplex the displays. We use the same set of pins to drive the anodes in every tube, while the grids select which tube we wish to control at any particular time.

The second question a lot of people ask is: when it's connected to +25V, why doesn't the grid 'suck up' all the electrons itself? The answer is that it actually does collect a few of them. However, the wires forming the grid are so fine that they occupy only about 15% of the available area, leaving around 85% open for the electrons to whizz through and strike the positively charged anodes beyond.

The final step is to add a phosphor coating to the anode, as shown in Fig.4(d). Later, when we add more anodes, each will receive its own phosphor coating. A phosphor is simply a fluorescent material.

Fluorescence means that the material emits visible light when it's excited by an external energy source. In our case, that energy source is a stream of electrons emitted by the heated filament. When a particular anode is driven to +25V, electrons strike its phosphor coating, causing it to fluoresce and produce the familiar blue-green glow of a VFD.

Once all of these components have been assembled inside the glass envelope, the air is evacuated to create a high vacuum, after which the tube is sealed. Without this vacuum, the electrons would constantly collide with air molecules, and very few would ever reach the anodes.

Learning our ABCs

Photo 2 shows a close-up of an inactive 7-segment VFD. You can see the unlit segments in creamy white, and the grid is clearly visible. Spotting the filament is a bit trickier, but it's there when you look closely.

The next point is that we need some way to identify the segments so we can talk about them meaningfully. By convention, the seven segments are labelled **a** through **g** (Fig.5). I'm not

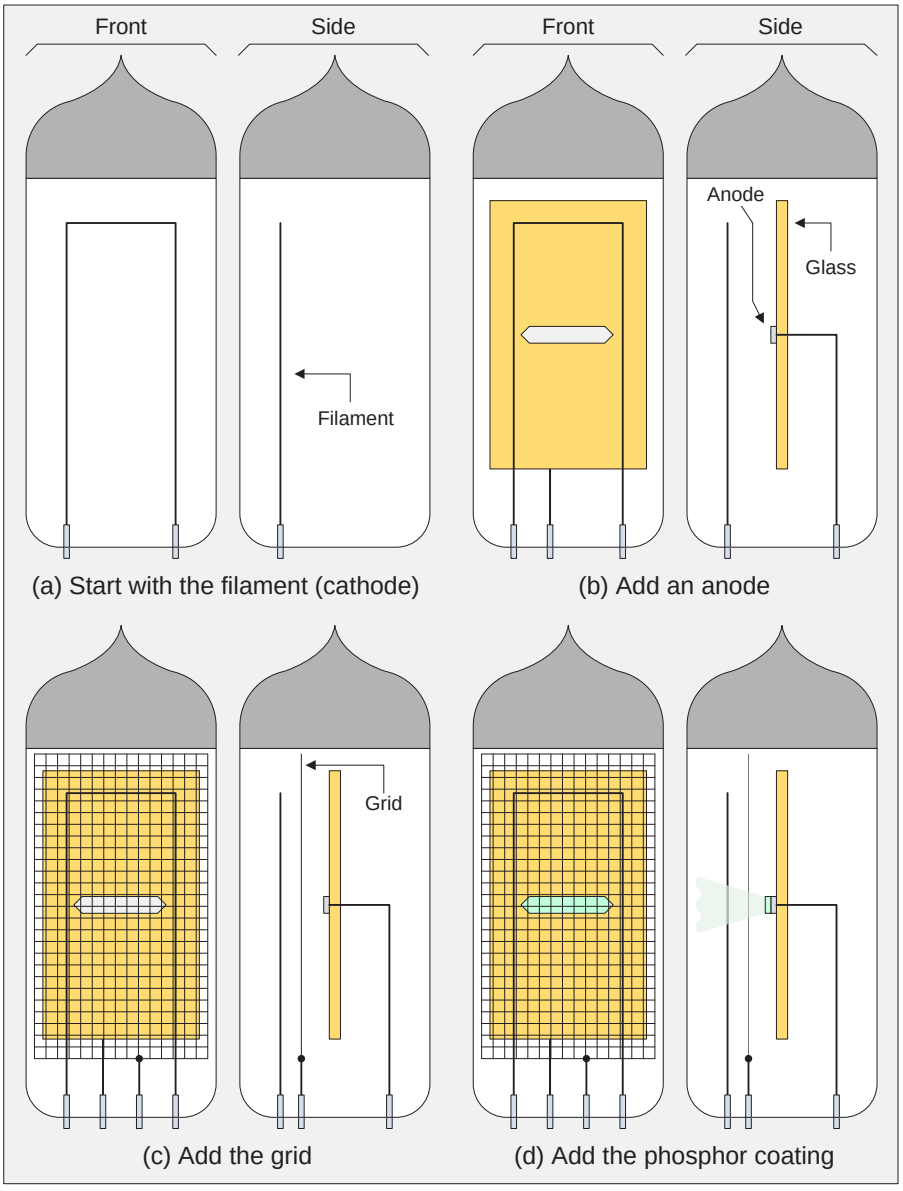

Fig.4: how VFD display tubes are made.

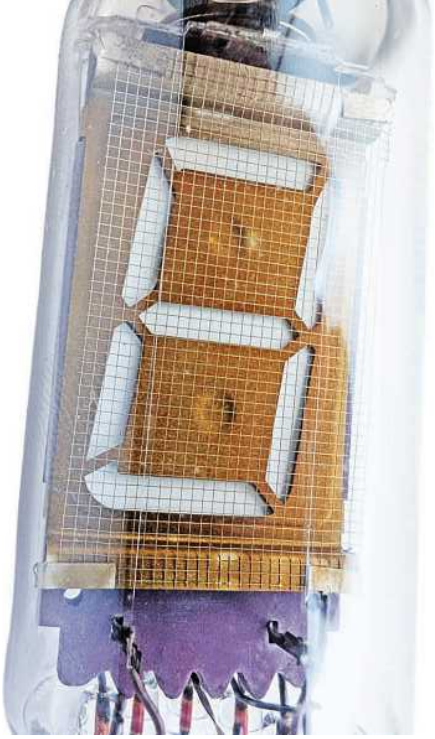

Photo 2: a VFD 7-segment display. Source: OSKAR (https://display-tubes.org/vfd/reflector-iv-12).

sure where this convention originated, but it became firmly established with the advent of commercial seven-segment displays and decoder integrated circuits (ICs) in the late 1960s and early 1970s.

In the case of displays that also sport a decimal point segment, the modern convention is to call it **dp** (for "decimal point", obviously). You do sometimes see it called **h** on older VFD and LED reference materials.

The **a-h** notation was fairly common in early Japanese and European datasheets, and you'll occasionally see it in Soviet-era and Eastern European documentation as well. The logic is simply that there are eight independently controllable elements, so they are labelled **a** through **h**.

We're using IV-11 VFDs. These are nice, medium-sized tubes—large enough to be easy to read, but small enough to be practical—and they're still relatively common. Although they are no longer being manufactured, large quantities of unused Soviet-era stock remain in circulation. These 'new old stock' (NOS) tubes are readily available from specialist suppliers and places like eBay.

These tubes have 12 pin positions but only 11 pins. Looking at the bottom of the tube, as shown in Fig.6(a), the pins are numbered clockwise, starting with pin 1 in the first position immediately after the empty location.

One thing we know for sure (culled from all the data sheets and scribbled notes I've tracked down on the interweb) is that pins 1 and 11 are the heater pins, to which we will be

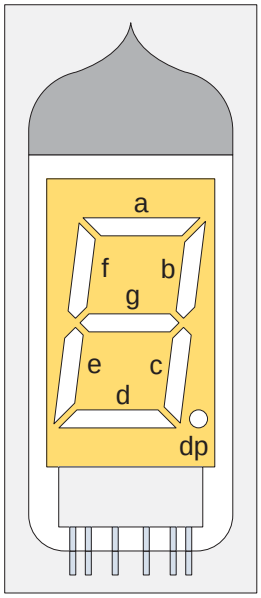

Fig.5: the standard 7-segment labels.

applying 1.5V DC. Another thing we know for sure (the same qualification applies) is that pin 2 is the grid, to which we will be applying +25V.

This means that pins 3 through 10 must be connected to the anodes **a-g** and **dp**, but which pin is connected to which anode?

I'm glad you asked. Different sources give different answers, so we are obliged to fall back on the time-honoured engineering practice of 'suck it and see' (ie, power things up and apply a voltage to each pin in turn to discover which segment lights).

Testing, testing, one two...

What we're going to do now is build a little test rig. For this, we'll need a breadboard (a half-size one will do), a single-pole, 12-throw (ie, 12-way) rotary switch, a 14-pin zero insertion

Photo 3: a 14-pin ZIF socket (top and bottom).

force (ZIF) socket, some header pins and some flying leads.

With the exception of the rotary switch and ZIF socket, you should already have everything else from our previous projects, but I've provided sources for everything at the end of this column, just in case.

A ZIF socket is an awesome little device. As seen in Photo 3, this is a special type of IC socket with a small locking lever. Lift the lever and a device can drop into place with virtually no force. Push the lever down, and the device's pins are gripped firmly by the socket contacts. We'll use this to test our VFD tubes without risking damage to their delicate pins.

Speaking of pins, the ZIF socket follows the same pin-numbering convention as for a standard dual in-line package (DIP) IC. Looking down on the socket from above, pin 1 is identified by the notch (or the corner with

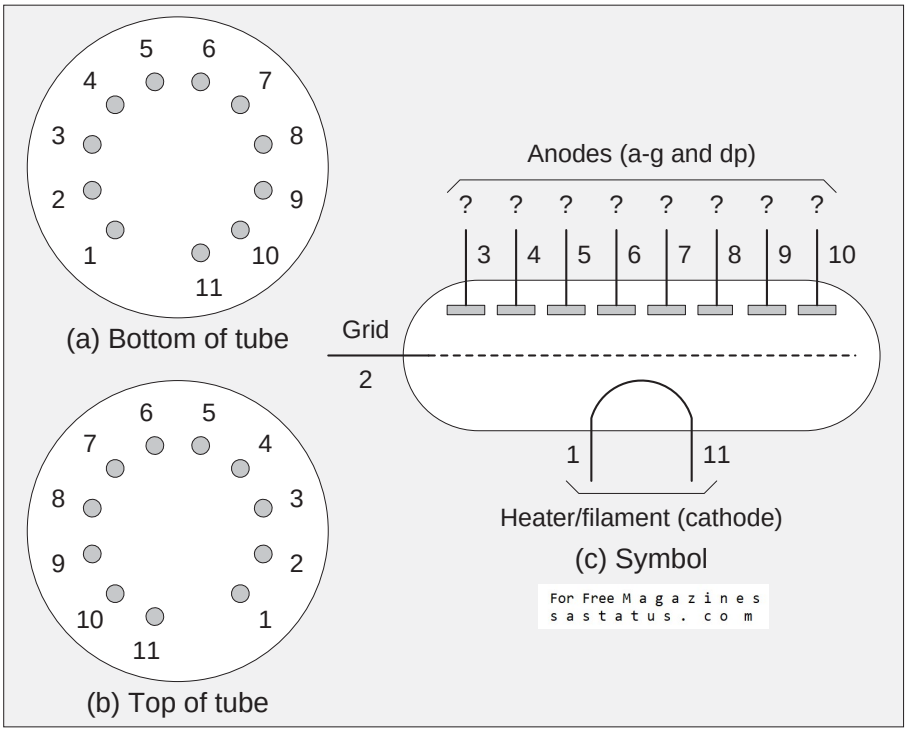

Fig.6: this shows how we will number the VFD tube pins.

the lever, in this case), and the pins are numbered anticlockwise around the package.

The pins on the base of the ZIF socket are too short to make reliable connections when plugged into a breadboard (these sockets are usually soldered directly to printed circuit boards). That's why we need the 0.1-inch/2.54mm pitch header pins. Cut two 7-pin strips and solder their short ends to the ZIF socket's pins. The long ends will then plug securely into the breadboard.

In the case of our single-pole, 12-throw rotary switch, the pole terminal is in the centre with the 12 throw terminals arranged in a circle around the periphery, as shown in Photo 4.

Unfortunately, there is no universal standard for numbering rotary switch terminals. The most common conventions are:

- Numbered clockwise when viewed from the shaft (front). This is fairly common in older switches and some military/aerospace documentation.
- Numbered clockwise when viewed from the solder terminals (rear). This is also very common, especially in datasheets, because that's the view you're looking at when wiring the switch.
- No numbering at all. Many manufacturers simply label the common terminal **C**, **COM** or **P**, and identify the switched terminals only in a drawing.

The problem is that 'clockwise' depends on which side you're looking from. A switch viewed from the rear appears to run anti-clockwise when viewed from the front.

For the purpose of these discussions, we're going to number the switch terminals the way I think of them in my head (which is the best place to think of them, in my experience): clockwise when viewed from the shaft (front), as illustrated in Fig.7(a).

Photo 4: a standard rotary switch.

When we eventually get to perform our tests, we'll start by rotating the shaft fully anti-clockwise to the pin 1 position. At that point, we'll want everything to be off, so we'll leave pin 1 unconnected, and then connect pins 2 through 9 on the switch (via the ZIF socket) to pins 3 through 10 on the VFD, respectively.

Let's start by taking one of our flying leads. It doesn't matter which type (male-male or male-female). Keep a male pin on one end, cut the other end off, bare the wire and solder it to the switch's pole pin. Repeat for switch pins 2 through 9, then wire everything up as illustrated in Fig.8.

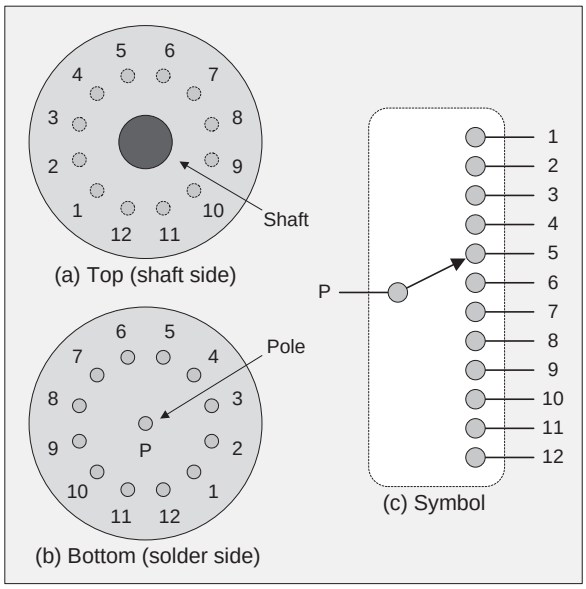

Fig.7: rotary switch pin numbering.

The numbers shown on the ZIF socket graphic are the pin numbers associated with the socket. Observe that the pin marked with a dot on the ZIF socket (the ZIF's pin 10) is the one into which we will plug pin 1 from the VFD. The remaining pins will relate as follows:

- ZIF.1 = VFD.6
- ZIF.2 = VFD.7
- ZIF.3 = VFD.8
- ZIF.4 = VFD.9
- ZIF.5 = VFD.10
- ZIF.6 = VFD.11 (0V)
- ZIF.7 = Unused
- ZIF.8 = Unused
- ZIF.9 = Unused
- ZIF.10 = VFD.1 (+1.5V)
- ZIF.11 = VFD.2 (Grid, +25V)
- ZIF.12 = VFD.3
- ZIF.13 = VFD.4
- ZIF.14 = VFD.5

Next time

I'm going to give you some time to acquire the bits and pieces you'll need to implement our VFD test setup shown in Fig.8. In our next column, we'll test our VFDs, then solder them onto our Chronological Contraption circuit boards.

As always, if you have any thoughts you'd care to share on anything you've read here, please feel free to email me at max@clivemaxfield.com PE

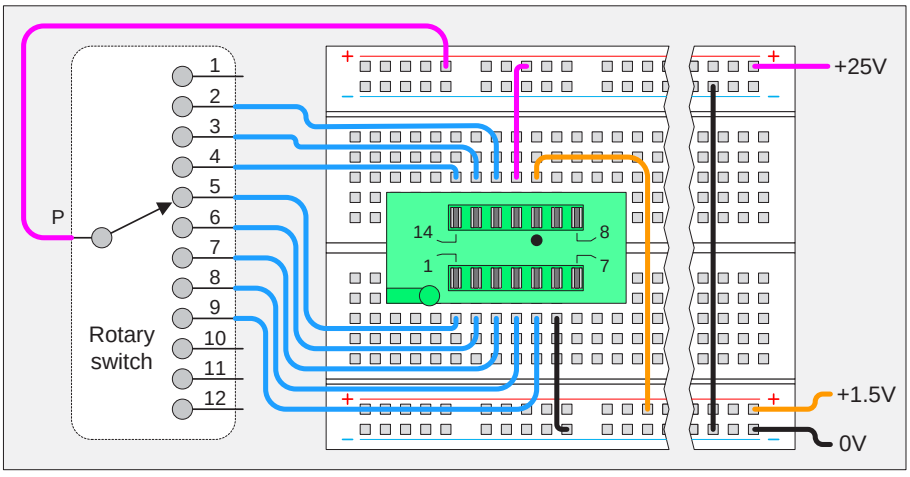

Fig.8: the test setup to verify our VFD tubes.

Devices and components used in this issue

Qty	Part	Source
1	1-pole 12-way rotary switch	www.amazon.co.uk/dp/B0CN8VL1BC
1	14-pin ZIF socket	www.amazon.co.uk/dp/B0C2VDCKXV
1	Half-size breadboard	www.amazon.co.uk/dp/B0CZ41KZ2H
14	Header pins (0.1-inch/2.54mm pitch)	www.amazon.co.uk/dp/B0DZT64TJJ
20	Flying leads	www.amazon.co.uk/dp/B01EV70C78

Practical Electronics PCB SERVICE

PROJECT	CODE	PRICE
OCTOBER 2026		
Ducted Heat Transfer Controller main PCB	17101251	£6.99
Ducted Heat Transfer Controller control panel PCB	17101252	£3.99
Ducted Heat Transfer Controller sensor PCB	17101253	£2.99
Model Railway DCC Decoder PCB	09111241	£2.99
SEPTEMBER 2026		
Home Automation Satellite PCB	15104251	£2.99
Mic the Mouse PCB set	08105251-2	£3.99
AUGUST 2026		
RP2350B Computer (assembled)	07204251A	£54.99
USB-C Power Monitor PCB set (including flex cable)	04102251-3	£6.99
Audio Out Phase-Linear Crossover PCB	AO-AUG-26	£4.99
VFD Chronometer digit board	CBAUG26-1	£2.99
VFD Chronometer spacer board	CBAUG26-2	£2.99
JULY 2026		
RP2350B Development Board (assembled)	07107251A	£24.99
SmartProbe PCB	P9054-04	£4.99
SWD Adaptor board for the SmartProbe	P9045-A	£2.99
Audio Out New All-Pass Filter PCB	AOJUL26-1	£4.99
Audio Out Alps Polarity Switch PCB	AOJUL26-2	£2.99
Audio Out E-switch Polarity Switch PCB	AOJUL26-3	£2.99
JUNE 2026		
RGB LED 'Analog' Clock PCB (black)	19101251	£7.99
USB Power Adaptor PCB (set of two)	18101251	£2.99
433MHz Receiver Module PCB	15103252	£2.99
Audio Out All-Pass Filter PCB set (three PCBs)	AO-JUN-26	£6.99
MAY 2026		
SSB Shortwave Receiver three PCB set	CSE250202-4	£15
Tool Safety Timer PCB	10104251	£3.99
433MHz Transmitter Module PCB	15103251	£2.99
APRIL 2026		
Versatile Battery Checker PCB	11104251	£3.99
Versatile Battery Checker front panel (black, 0.8mm)	11104252	£4.99
Power Rail Probe PCB	P9058-1-C	£3.99
Pico 2 Audio Analyser PCB/panel (black)	04107231	£3.99
MARCH 2026		
USB-Programmable Frequency Divider PCB (blue)	04108241	£3.99
Rotating Light PCB (black)	09101251	£2.99
433MHz Transmitter PCB	15103251	£2.99
Arduino Uno Game Cartridge PCB (Cool Beans)	CBmar26uno	£3.99
Arduino Nano Game Cartridge PCB (Cool Beans)	CBmar26nano	£2.99
Silent Stomp Switch PCB pair (Audio Out)	AO-MAR26-*	£3.99
FEBRUARY 2026		
Power LCR Meter PCB	04103251	£7.99
Versatile Waveform Generator PCB	04104251	£5.99
JANUARY 2026		
High-Bandwidth Differential Probe	9015-D	£5.99
NFC-Programmable IR Remote Keyfob	15109231	£2.99
Pre-driver Overdrive Pedal main PCB	AO-JAN26-1	£4.99
Updated jack PCB for the Pre-driver	AO-JAN26-2	£3.99
DECEMBER 2025		
Variable Speed Drive (black)	11111241	£11.99
Digital Capacitance Meter	04111241	£5.99
Battery-Powered Model Train transmitter	09110241	£3.99
Battery-Powered Model Train through-hole receiver	09110242	£3.99
Battery-Powered Model Train SMD receiver	09110243	£3.99
Battery-Powered Model Train battery charger	09110244	£3.99
Transmanium Fuzz Box jack PCB	AO-DEC-25	£3.99

PROJECT	CODE	PRICE
NOVEMBER 2025		
Capacitor Discharger	9047-01	£4.99
5MHz, 40A Current Probe	9049-01	£5.99
Transmanium Fuzz Box main PCB	AO-NOV-25	£4.99
OCTOBER 2025		
Seaside Sound Simulator PCB (blue)	01111241	£7.99
Fender Jazz Bass pickguard	23109241	£5.99
J&D Luthiers T-Style pickguard	23109242	£5.99
Music Man Stingray pickguard	23109243	£5.99
Fender Telecaster pickguard	23109244	£4.99
3D Printer Filament Dryer control PCB	28110241	£4.99
SEPTEMBER 2025		
Automatic LQ Meter PCB	CSE240203A	£5.99
Automatic LQ Meter front panel (black)	CSE240204A	£5.99
Compact Hi-Fi Headphone Amplifier PCB (blue)	01103241	£4.99
Discrete Ideal Bridge Rectifier TH or SMD PCB	18108241/2	£2.99
Dual Rail Supply Protector PCB	18109241	£3.99
Audio Out Cascode Amplifier PCB	CASCODEA	£4.99
Audio Out Cascode Amplifier Power Supply PCB	CASCODEP	£4.99
AUGUST 2025		
Micromite Explore-40 PCB	07106241	£2.99
Micromite Explore-40 Audio Breakout PCB	07101222	£2.99
8-Channel Learning IR Remote Receiver PCB	15108241	£5.99
FlexiDice PCB set	08107241-2	£4.99
JULY 2025		
Compact OLED Clock & Timer PCB	19101231	£4.99
180-230V DC Motor Speed Controller PCB	11104241	£9.99
Mains Power-Up Sequencer PCB (same as Feb '25)	10108231	£11.99
JUNE 2025		
Styloclone (case version)	23106241	£6.99
Styloclone (standalone version)	23106242	£7.99
Dual Mini LED Dice (TH LEDs, white solder mask)	08103241W	£2.99
Dual Mini LED Dice (TH LEDs, black solder mask)	08103241B	£2.99
Dual Mini LED Dice (SMD LEDs, white solder mask)	08103242W	£2.99
Dual Mini LED Dice (SMD LEDs, black solder mask)	08103242B	£2.99
Dual Hybrid Supply regulator board	18107211	£5.99
Dual Hybrid Supply front panel control board	18107212	£2.99
Dual Hybrid Supply CPU board	01106193	£3.99
Dual Hybrid Supply LCD adaptor	01106196	£2.99
MAY 2025		
Skill Tester 9000	08101241	£9.99
USB-C Serial Adaptor (black solder mask)	24106241	£2.99
DC Supply Protector (SMD adjustable version)	08106241	£3.99
DC Supply Protector (TH adjustable version)	08106242	£3.99
DC Supply Protector (TH fixed version)	08106243	£3.99
APRIL 2025		
ESR Test Tweezers PCB set	04105241+	£5.99
Fan Speed Controller Mk2	10104241	£4.99
Calibrated Microphone (SMD version)	01108231	£3.99
Calibrated Microphone (TH version)	01108232	£3.99
MARCH 2025		
Pico Gamer	08104241	£9.99
Compact Frequency Divider	04112231	£4.99
FEBRUARY 2025		
Laser Communicator transmitter	16102241	£3.99
Laser Communicator receiver	16102242	£3.99
Model Railway Points Controller transmitter	09101241	£3.99
Model Railway Points Controller receiver	09101242	£2.99
Mains Power-Up Sequencer	10108231	£11.99

PCBs for most recent *PE/EPE* constructional projects are available. From the July 2013 issue onwards, PCBs with eight-digit codes have silk screen overlays and, where applicable, are double-sided, have plated-through holes, and solder mask. They are similar to photos in the project articles. Earlier PCBs are likely to be more basic and may not include silk screen overlay, be single-sided, lack plated-through holes and solder mask.

Always check price and availability in the latest issue or online. A large number of older boards are listed for ordering on our website.

In most cases we do not supply kits or components for our projects. For older projects it is important to check the availability of all components before purchasing PCBs.

Magazine back issues are available – see https://www.electronpublishing.com/product-category/pe-back-issues-print-and-digital/

Double-sided | plated-through holes | solder mask

PROJECT	CODE	PRICE
JANUARY 2025		
USB to PS/2 Converter for VGA PicoMite	07111231	£3.99
ps2x2pico USB to PS/2 Converter	07111232	£3.99
Compact Microphone Preamplifier (case version)	01110231	£5.99
Compact Microphone Preamplifier (integrated version)	01110232	£5.99
DECEMBER 2024		
Raspberry Pi Clock Radio PCB pair	19101241/2	£12.99
VGA PicoMite (blue)	07107221	£5.99
Secure Remote Switch receiver	10109231	£4.99
Secure Remote Switch transmitter (module version)	10109232	£2.99
Secure Remote Switch transmitter (discrete version)	10109233	£2.99
NOVEMBER 2024		
Ideal Diode Bridge Rectifiers (28mm square)	18101241	£3.99
Ideal Diode Bridge Rectifiers (22mm square)	18101242	£3.99
Ideal Diode Bridge Rectifiers (4-pin SIL)	18101243	£3.99
Ideal Diode Bridge Rectifiers (tiny SOT-23)	18101244	£3.99
Ideal Diode Bridge Rectifiers (standalone SMD)	18101245	£4.99
Ideal Diode Bridge Rectifiers (standalone TH; 2oz copper)	18101246	£6.99
Multi-channel Volume Control (volume PCB)	01111221	£4.99
Multi-channel Volume Control (control PCB)	01111222	£4.99
Multi-channel Volume Control (OLED PCB, black)	01111223	£4.99
Coin Cell Emulator (black)	18101231	£3.99
OCTOBER 2024		
TQFP-32 Programming Adaptor	24108231	£4.99
TQFP-44 Programming Adaptor	24108232	£4.99
TQFP-48 Programming Adaptor	24108233	£4.99
TQFP-64 Programming Adaptor	24108234	£4.99
2m VHF FM Test Signal Generator	06107231	£4.99
SEPTEMBER 2024		
PIC Programming Adaptor	24105231	£4.99
Salad Bowl Speaker crossover	01109231	£7.99
AUGUST 2024		
Arduino ESR Meter combined L/C-ESR	04106182	£9.99
Arduino ESR Meter separate ESR	04106183	£7.99
Watering System Controller	15110231	£11.99
JULY 2024		
Dynamic NFC Tag (22 × 31mm)	06101231	£5.95
Dynamic NFC Tag (86 × 54mm)	06101232	£7.95
Reciprocal Frequency Counter	CSE230101C	£8.95
Reciprocal Frequency Counter front panel (black mask)	CSE230102	£7.95
Pi Pico-based Thermal Camera	04105231	£6.95
JUNE 2024		
Loudspeaker Test Jig	04106231	£12.95
MAY 2024		
GPS-Disciplined Oscillator	04103231	£12.95
Dual RF Amplifier	CSE220602A	£5.95
Songbird	08103231	£7.95
APRIL 2024		
Basic RF Signal Generator	CSE221001	£8.99
Basic RF Signal Generator front panel (black solder mask)	CSE220902B	£9.99
Silicon Chirp Cricket	08101231	£7.99
Test Bench 'Swiss Army Knife'	04110221	£9.99
MARCH 2024		
Digital Volume Control Potentiometer SMD version	01101231	£8.95
Digital Volume Control Potentiometer through-hole version	01101232	£9.95
Advanced SMD Test Tweezers body	04106221	£6.95
Advanced SMD Test Tweezers legs (pair)	04106212	£7.95

PROJECT	CODE	PRICE
FEBRUARY 2024		
Active Mains Soft Starter	10110221	£10.95
JANUARY 2024		
Q Meter	CSE220701	£8.95
Q Meter front panel (black solder mask)	CSE220704	£7.95
Raspberry Pi Pico W BackPack	07101221	£8.95
DECEMBER 2023		
Digital Boost Regulator	24110224	£8.95
Dual-Channel Power Supply for Breadboards	04112221	£11.95
Display Adaptor for the Breadboard PSU	04112222	£9.95
NOVEMBER 2023		
LC Meter Mk3	CSE220503C	£9.95
DC Supply Filter for vehicles	08108221	£8.95
Discrete Audio Op Amp PCB	AO1-JUL23	£9.95
Discrete Audio Op Amp PCB + essential components	AO2-JUL23	£17.95
OCTOBER 2023		
Buck/Boost Charger Adaptor	14108221	£9.95
PIC Breakout Board for SOIC parts	24110225	£5.95
PIC Breakout Board for DIP parts	24110222	£5.95
AVR64DD32 Breakout board	24110223	£5.95
Automatic Train Controller	09109221	£6.95
Chuff Sound module	09109222	£6.95

For the many earlier PCBs that we stock please see the PE website: www.electronpublishing.com

PE/EPE PCB SERVICE

Order Code Project Quantity Price

. . . . . .

. . . . . .

. . . . . .

. . . . . .

. . . . . .

. . . . . .

. . . . . .

Name . . . . . .

Address . . . . . .

. . . . . .

Tel . . . . . .

Email . . . . . .

Card No . . . . . .

Valid From . . . . . Expiry Date . . . . . .

Card Security No . . . . . MasterCard VISA

You can also order PCBs by phone, email or via the shop on our website: www.electronpublishing.com

No need to cut your issue – a copy of this form is just as good!

All prices include VAT and UK p&p. Add £4 per *project* for post to Europe; £5 per *project* outside Europe.

Orders and payment should be sent to:

Practical Electronics, Electron Publishing Ltd
113 Lynwood Drive, Merley, Wimborne, Dorset BH21 1UU
Tel 07918 614662 Email: pe@pemag.au
Online Shop: www.electronpublishing.com

Sorry, cheques cannot be accepted as we are based in Australia and our UK bank account does not allow cheque deposits.

NOTE: Most boards are in stock and sent within seven days of order receipt. Please allow up to 28 days delivery if we need to restock.

CLASSIFIED ADVERTISING

Practical Electronics

If you want your advertisements to be seen by the largest readership at the most economical price then our classified page offers excellent value.

The rate for semi-display space is £10 per centimetre high, with a minimum height of 2·5cm. All semi-display advertisements have a width of 5.5cm. The prepaid rate for classified adverts is 40p per word (minimum 12 words).

For rates, further information on display and classified advertising or to book an ad, please email us at **pe@pemag.au**

Practical Electronics reaches more UK readers than any other UK monthly hobby electronics magazine. Our sales figures prove it.

We have been the leading monthly magazine in this market for more than thirty years.

W5SWL Electronics

Technical Products and Small Component Parts, RF and Audio Connectors, Hobby Electronics, Ham Radio to Microwave Products, General Hobby Supplies, Technical Books, Tools, and More. **Located in the United States.**

Visit **www.W5SWL.com**

BOWOOD ELECTRONICS LTD

Suppliers of Electronic Components

www.bowood-electronics.co.uk

Unit 10, Boythorpe Business Park, Dock Walk, Chesterfield, Derbyshire S40 2QR. Sales: 01246 200 222

Send large letter stamp for Catalogue

COAST ELECTRONICS

BREAKOUTS-COMPONENTS-CONTRACT DESIGN-3D PRINTER PARTS-MUSICAL-MICROCONTROLLERS

WWW.COASTELECTRONICS.CO.UK

MISCELLANEOUS

VALVES AND ALLIED COMPONENTS? For free stock list and/or advice, please contact me: geoffdavies337@gmail.com Telephone: 01788 574774

PIC DEVELOPMENT KITS, DTMF kits and modules, CTCSS Encoder and Decoder/Display kits. Visit www.cstech.co.uk

Electrical Industries Charity (EIC)
We help people working in the electrical, electronics and energy community as well as their family members and retirees.

We use workplace programmes that give the industry access to financial grants and a comprehensive range of free and confidential services.

www.electricalcharity.org

ADVERTISING INDEX

AOSHOP 55
ENIGMA COMPONENT SHOP. 7
HAMMOND MANUFACTURING. 5
MICROCHIP TECHNOLOGY . . . . . Outer back cover
PCBWAY. . . . . Inner front cover
PDF BACK ISSUE DOWNLOADS 10, 22
PO LABS 37
SILICON CHIP MAGAZINE. 23

Advertising mailing address

Electron Publishing Australia
PO Box 194
Matraville NSW 2036
Australia

Email pe@pemag.au

Web www.electronpublishing.com

For editorial contact details, see page 2.

Coming Up – in the Nov/Dec issues

DIgital Preamplifier & Crossover

A comprehensive preamplifier with four stereo analog inputs, two digital inputs and USB audio in/out. It uses high-performance digital signal processing (DSP) for all functions.

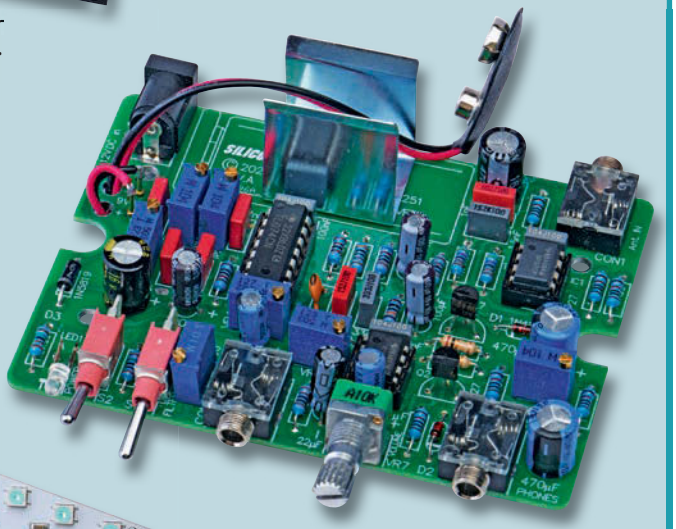

Earth Radio

Listen to the sounds of auroras, lightning and other atmospheric phenomena with this radio and a simple loop antenna.

RGB LED Christmas Star

- Easy assembly
- Over a dozen impressive effects
- Onboard switch-mode power supply

November issue on sale: 6th of October, 2026

Content may be subject to change

Practical Electronics

Subscriptions hotline

To subscribe or renew online, paying by credit/debit card or PayPal, please visit our shop at: www.electronpublishing.com

For those wishing to pay by direct debit, cheque or over the telephone, please call 'Select Publisher Services' on 01202 087631, or send a cheque (payable to: 'Select Publisher Services Ltd') with your details to:

Practical Electronics Subscriptions, PO Box 6337, Bournemouth BH1 9EH, United Kingdom

Remember, we print the date of the last issue of your current subscription in a box on the address sheet that comes with your copy.

Digital subscribers, please call 07918 614662 or visit: www.electronpublishing.com

Published on approximately the first Thursday of each month by Silicon Chip Publications, PO Box 194 Matraville NSW 2036 Australia. Printed in England by The Manson Group, Reynolds House, 8 Porters Wood, Valley Road Industrial Estate, St Albans AL3 6PZ. Distributed by Seymour, 86 Newman St, London W1T 3EX. Subscriptions UK: £33.99 (6mo); £59.99 (12mo.); £114.99 (2yrs). EUROPE: airmail service, £38.99 (6mo); £69.99 (12mo); £129.99 (2yrs). REST OF THE WORLD: airmail service, £46.99 (6mo); £84.99 (12mo); £164.99 (2yrs). Make payments to 'Practical Electronics', Practical Electronics Subscriptions, PO Box 6337, Bournemouth BH1 9EH United Kingdom. Email: pesubs@selectps.com. PRACTICAL ELECTRONICS is sold subject to the conditions that it shall not, without the written consent of the Publishers first having been given, be lent, resold, hired out or otherwise disposed of by way of Trade at more than the recommended selling price on the cover, and that it shall not be lent, resold, hired out or otherwise disposed of in a mutilated condition or in any unauthorised cover by way of Trade or affixed to or as part of any publication or advertising, literary or pictorial matter.

Did you know our online shop now sells the current issue of *PE* for £5.99 inc. p&p?

Practical Electronics
The UK's premier electronics, computing and maker magazine
Max's Cool Beans
Cellular automata: The Game of Life
Net Work
Uses for 18V/21V Li-ion tool batteries
Audio Out
RetroTech report part 2
Circuit Surgery
Accurately simulating 555s with LTspice
DIY DCC Model Railway Part 1: DCC Decoder
Ducted Heat Transfer Controller
Rigol DHO924S Digital Storage Oscilloscope Review
PLUS!
Techno Talk – new web development & testing techniques
The Fox Report – Humanoid robots are just around the corner
Advanced uses for HomeAssistant running on a Raspberry Pi
Oct 2026 £5.99
www.electronpublishing.com
@practicalelec
practicalelectronics
Practical Electronics
Audio Out
Visiting Retro-Tech 2026
Circuit Surgery
Measuring resistance with a 555 timer chip
Mic the Mouse:
a sound-sensitive pet that won't make a mess
Sep 2026 £5.99
www.electronpublishing.com
@practicalelec
practicalelectronics

You read that right! We now sell the current issue of your favourite electronics magazine for exactly the same price as in the High Street, but we deliver it straight to your door – and for UK addresses we pay the postage. No need to journey into town to queue outside the newsagent. Just go to our website, set up an account in 30 seconds, order your magazine and we'll do the rest.

www.electronpublishing.com

Empower Your IoT Interfaces

Wireless Connectivity Meets Advanced Touch and Graphics With PIC32-BZ6

Build smarter and faster with one chip.

The PIC32-BZ6 MCU family simplifies connected product design by integrating wireless, touch, graphics, motor control and industrial interfaces into a single, high-performance IC. Ideal for smart home, industrial, medical and automotive devices, it reduces component count, accelerates development and future-proofs your designs with support for Bluetooth®, Thread® and Matter.

Key Features

- One chip: wireless + UI + real-time control
- Bluetooth LE and 802.15.4 (Thread, Matter)
- Integrated capacitive touch and graphics display controller
- Industrial-grade interfaces: CAN-FD, USB, Ethernet
- 2 MB Flash and 512 KB RAM for complex applications and OTA
- Simplify your architecture, cut development time and deliver powerful connected experiences all with a single device

MICROCHIP

microchip.com/bluetooth-hmi

The Microchip name and logo and the Microchip logo are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries. All other trademarks are the property of their registered owners.
© 2026 Microchip Technology Inc. All rights reserved.
MEC2657A-UK-07-26